



User Guide

Amazon Quick



Amazon Quick: User Guide

Copyright © 2026 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

Amazon's trademarks and trade dress may not be used in connection with any product or service that is not Amazon's, in any manner that is likely to cause confusion among customers, or in any manner that disparages or discredits Amazon. All other trademarks not owned by Amazon are the property of their respective owners, who may or may not be affiliated with, connected to, or sponsored by Amazon.

Table of Contents

What is Amazon Quick?	1
Benefits of Quick	1
Accessing Quick	2
Getting started	2
How Amazon Quick works	3
You start in chat	3
Agents process your requests	3
Spaces organize context	3
Integrations connect Quick to your world	3
Features use those building blocks	4
Where to go next	5
Setting up and signing in	6
Complete initial configuration tasks	7
Sign up for an AWS account	7
Signing up at aws.com/quick	7
Who is this for?	8
Sign-in options	8
Step-by-step signup procedure	8
What happens behind the scenes	12
Signing up through the AWS Console	12
Non-admin sign-up	13
Signing in to Quick	14
Signing in to a Free or Plus account	15
Post-setup configuration	15
Signing in through the AWS Console	16
Signing in with multiple accounts	18
Getting started with Amazon Quick	22
Using the Amazon Quick console	22
Using the Amazon Quick menu and landing page	22
Searching Amazon Quick	26
Choosing a language in Amazon Quick	26
Getting started with Amazon Quick Sight	28
Quick start: Create an analysis using sample data	28
Create a dashboard using sample data	29

Sample data	35
Administering Amazon Quick	37
Free, Plus, and Max accounts	37
Plans and pricing	38
Upgrading your plan	39
Managing your account	44
Managing users	47
Managing assets	51
Plan and billing	54
Understanding usage	58
Account settings	62
User profile and preferences	65
Enterprise accounts	70
Getting started	72
Setting up	89
Configuring VPC connections	164
Configuring Amazon Quick Sight access to AWS data	185
Managing account	230
Managing subscriptions	237
Managing identities	254
Managing security	276
Managing permissions	317
Governance	401
Managing customizations	419
Amazon Quick usage metrics	430
AWS security	437
Ask questions and explore data using Amazon Quick chat	532
Welcome message	533
Suggested prompts	533
Expand and collapse modes	533
Agent picker	533
Scoping response generation to specific data	534
Chatting with datasets	534
Chatting with Topics	535
Web search	536
Upload files and chat	536

Actions	537
Flows	537
Memory and response personalization	538
Source citations	539
Copy responses	539
Provide feedback	539
Artifact panel	540
Conversation management	540
Response events	541
Explanations	541
Contextual awareness	541
Managing your preferences in Amazon Quick	542
Customizing your chat experience	542
Setting your default chat agent	542
Personalizing your experience	543
Managing your memories	543
Accessing User Preferences	543
Chat explanations	543
Chat with your dashboards	544
Chat with your datasets	547
Document and visual creation	549
Supported document and visual types	549
How document and visual creation works	550
Creating documents and visuals in chat	551
Example prompts	551
Editing existing documents	553
Working with uploaded files	555
Using templates	555
PowerPoint templates	556
Excel templates	556
Word templates	557
Previewing and downloading	557
Business visual creation	558
Standalone infographics and visuals	558
Visuals embedded in documents	559
Brand theming	559

Best practices	560
Limitations	561
Create and deploy AI assistants using Amazon Quick chat agents	562
Amazon Quick chat agent types	563
User interaction with chat agents	564
User permissions for chat agents	565
Custom permissions for chat agents	566
System chat agent	568
Custom chat agents	569
Create custom chat agent	569
Share custom chat agent	575
Embed custom chat agents	576
Update custom chat agent	577
Manager access to a chat agent	578
View chat agents	579
View chat agent details	579
Duplicate custom chat agents	580
Delete chat agents	580
Custom chat agent context configuration	581
Agent context types	581
Use a chat agent	584
Organize and collaborate using Amazon Quick spaces	586
User interactions with spaces	587
User permissions for spaces	587
File upload status	588
Limitations of spaces	589
Creating spaces	590
Managing spaces	590
Viewing spaces	591
Editing spaces	591
Adding space knowledge	592
Viewing content in a space	592
Using spaces	593
Sharing spaces	594
Deleting spaces	594
Removing knowledge	595

Adding external data sources	595
Adding application actions	596
Visualize, analyze, and share data with Amazon Quick Sight	597
Key concepts	598
Connecting to data	599
Supported data sources	601
Connect to your data	606
Data catalog integration (Preview)	655
Data source quotas	671
Supported data types and values	673
Working with datasets	686
Working with data sources	752
Refreshing data	757
Importing data into SPICE	759
Refreshing SPICE data	761
Using SPICE data in an analysis	766
View SPICE ingestion history	767
Troubleshooting skipped row errors	767
SPICE ingestion error codes	770
Updating files in a dataset	775
Preparing data	776
Data Preparation Experience (New)	777
Describing data	808
Choosing file upload settings	809
Data Preparation Experience (Legacy)	810
Using SQL to customize data	1113
Adding geospatial data	1116
Using unsupported or custom dates	1173
Adding a unique key to a Quick Sight dataset	1174
Integrating SageMaker AI models	1175
Preparing dataset examples	1186
Analyses and reports: Visualizing data	1193
Working with an analysis	1193
Adding and managing sheets	1217
Working with interactive sheets in Amazon Quick Sight	1218
Working with pixel perfect reports in Amazon Quick Sight	1261

Working with items on sheets	1276
Using themes in Amazon Quick Sight	1511
Keyboard shortcuts	1518
Dashboards and reports: Sharing data	1521
Publishing dashboards	1522
Sharing dashboards	1527
Use connectors in dashboards	1541
Share your view of a dashboard	1544
Sending reports	1544
Subscribing to reports	1549
Threshold alerts	1550
Print a dashboard or analysis	1559
Exporting as PDFs	1559
PDF Error codes	1560
Organizing assets into folders	1562
Exploring dashboards	1575
Interacting with dashboards	1577
Interacting with pixel perfect reports	1586
Subscribe to emails and alerts	1587
Reader generated reports	1589
Bookmarks	1592
Gaining insights with ML	1596
Understanding the ML algorithm	1597
Dataset requirements	1601
Adding insights	1602
Autonarratives	1605
Detecting outliers	1640
ML-powered forecasts and what-ifs	1652
Generative BI with Quick Sight	1655
Get started	1657
Augmenting Amazon Quick Sight insights with Amazon Q Business	1660
Authoring experience	1663
Executive summaries	1666
Authoring Q&A	1667
Manage topic permissions through dashboards	1668
Turn on the Dashboard Q&A experience in Amazon Quick Sight	1669

Q&A null support	1671
Improve Q&A accuracy with custom instructions	1674
Asking and answering questions of data with Generative BI	1679
Opting out of Generative BI	1682
Working with Topics	1685
Working with data stories in Amazon Quick Sight	1703
Working with scenarios in Amazon Quick Sight	1708
Troubleshooting	1714
Resolving Amazon Quick Sight issues and error messages	1714
Athena issues	1715
Data source connectivity issues	1725
Login issues	1737
Visual issues	1744
Developing with Amazon Quick Sight	1748
Required knowledge	1748
Available API operations for Amazon Quick Sight	1749
Terminology and concepts	1750
Developing with the Amazon Quick Sight APIs	1752
Events integration	1760
Embedded analytics	1781
Using Amazon Quick Research	2022
What you can do with Quick Research	2022
Starting new research	2022
Enter research objective	2023
Select research materials	2024
Review and finalize research plan	2027
View research report	2028
Updating research	2029
Making updates	2029
Research versions	2030
Summarizing the research	2030
Custom summary	2031
Executive summary	2031
General share out	2031
Sharing the research	2031
Download the research	2031

Sharing the research with stakeholders	2032
Automating research with flows	2033
When to use research in Quick Flows	2033
Getting started	2033
Using third party data for research	2033
Using FactSet data	2034
Using IDC data	2036
Using S&P Global Energy data	2037
Using S&P Global Market Intelligence data	2039
Using US Patents data	2041
Using PubMed data	2042
Admin controls	2042
Turn on/off web search	2042
Status definitions	2042
Using Amazon Quick automations	2044
Types of automations	2044
Choosing between Quick Flows and Quick Automate	2044
Getting started	2044
Using Amazon Quick Flows	2045
What are Flows?	2045
Why Flows?	2046
Getting started	2046
Components and features	2050
Creating and managing	2057
Running flows	2070
Managing Quick Flows	2074
Quick Flows limits	2076
Using Amazon Quick Automate	2078
Overview	2078
Key concepts	2079
Getting started	2080
Building automations	2085
Automation actions	2096
Testing automations	2173
Deploying automations	2179
Export and import automations	2192

Monitoring performance	2194
Human-in-the-loop task center	2201
Build web applications with apps in Quick	2205
Key capabilities	2205
How apps in Quick works	2205
Getting started	2206
Prerequisites	2206
Open the apps in Quick landing page	2206
Create an app	2206
Edit an app	2207
Manage versions	2207
Review app settings	2208
Publish an app	2208
Share an app	2208
Manage public app usage	2209
What you can build	2210
User interfaces	2210
Data storage	2210
AI inference	2210
Embedded visuals	2211
Embedded chat experiences	2211
Public-facing applications	2211
File handling	2212
Connecting to integrations	2212
Prerequisites	2212
Integration types	2213
Authentication models	2213
Adding an integration to your app	2213
How apps in Quick works with other Quick capabilities	2213
Integration tips	2214
Working with data	2214
Built-in app storage	2214
Exporting data	2216
Handling write approvals when AI inference is active	2216
Accessing user information	2216
Security and sandbox	2217

Authentication	2217
Authorization layers	2217
Integration consent model	2218
Sandbox restrictions	2218
Public app security	2219
Tips and best practices	2219
Limitations	2220
Regional availability	2220
Limits	2220
Building experience	2221
Sharing and access	2221
Portability	2221
Integration limitations	2222
Sandbox restrictions	2222
Troubleshooting	2222
Recovering from content policy violations	2222
Recovering from a broken app	2223
Fixing connector authentication issues	2223
Debugging complex apps	2223
File export not working	2223
Integrations	2224
Supported integrations	2225
Integration capability definitions	2228
Authentication method definitions	2229
Set up integrations in the console	2229
Choose integration options	2229
View setup process examples	2230
Integration management options	2231
Integration workflows	2232
Creating a knowledge base	2232
Creating a connector	2232
Managing existing integrations	2233
Knowledge bases	2234
How knowledge bases work	2235
General workflow	2236
Common configuration settings	2236

Best practices for managing ACLs in knowledge bases	2239
Sharing knowledge bases and data sources	2242
Sync reports and observability	2246
Monitor index storage usage	2251
Troubleshooting knowledge bases	2254
Connectors	2258
How connectors work	2258
Types of actions	2259
Prerequisites	2261
Tool permissions	2261
Supported connector types and available actions	2263
Connector APIs	2270
Authentication methods	2279
Action execution methods	2281
Monitoring and maintenance	2282
Best practices	2284
Troubleshooting	2286
Integration-specific guides	2288
Adobe Marketing Agent integration	2288
Airtable integration	2292
S3 integration	2298
Asana integration	2317
Atlassian Confluence Cloud integration	2321
Atlassian Jira Cloud integration	2338
AWS Agent Registry	2345
BambooHR integration	2351
Box integration	2354
Canva integration	2361
Cisco Webex Meetings	2364
Coding agents	2367
Cisco Webex Video Messaging	2367
Dropbox integration	2370
Dun & Bradstreet	2378
GitHub integration	2380
Gmail integration	2382
Google Analytics integration	2388

Google Calendar integration	2395
Google Chat	2401
Google Docs integration	2404
Google Drive integration	2410
Google Meet integration	2430
Google Sheets integration	2436
Google Slides integration	2442
HG Insights integration	2447
HubSpot integration	2451
HuggingFace integration	2454
Intercom integration	2456
Linear integration	2458
Local connectors	2459
Microsoft OneDrive integration	2461
Microsoft OneNote	2496
Microsoft Outlook integration	2501
Microsoft SharePoint integration	2510
Microsoft Teams integration	2556
MCP integration	2566
Monday.com integration	2578
Moodys GenAI Ready Data	2579
New Relic integration	2582
Notion integration	2587
OpenAPI integration	2588
PagerDuty integration	2595
PagerDuty Advance integration	2596
QuickBooks integration	2598
REST API integration	2604
Salesforce integration	2608
SAP integrations	2619
ServiceNow integration	2630
Shopify integration	2636
Slack integration	2638
Smartsheet integration	2648
Snowflake Cortex Agent	2651
Visier Agent	2653

Web Crawler integration	2656
WhatsApp integration	2668
Zapier integration	2670
Zendesk integration	2673
Zoom integration	2678
ZoomInfo integration	2684
Support escalation path	2687
Bring Your Own Amazon Q Business Index	2688
Overview of Amazon Q Business indexes in Amazon Quick	2689
Prerequisites	2690
Supported authentication methods	2690
Setting up permissions	2691
Creating knowledge bases from Amazon Q Business indexes	2692
Sharing Amazon Q Business index knowledge bases	2692
Using Amazon Q Business index knowledge bases	2693
Limitations	2694
Billing	2695
Feature comparison	2696
Troubleshooting	2696
Security best practices	2698
User types and capabilities	2698
Common use cases	2698
Interacting with your Amazon Q Business indexes	2699
Bring your own Amazon Bedrock managed knowledge base	2700
Prerequisites	2701
Setting up permissions	2701
Creating knowledge bases from Amazon Bedrock managed knowledge bases	2702
Access control list (ACL) support	2703
Billing	2703
Limitations	2703
Security best practices	2704
Monitoring and observability	2704
Troubleshooting	2707
Connect Quick to existing tools using extensions	2709
User interactions with extensions	2710
User capabilities and permissions	2711

Administrator responsibilities	2712
Author capabilities	2713
End user access	2714
Extension access troubleshooting	2716
Amazon Quick browser extension	2716
Supported browsers	2717
Browser extension benefits	2717
Browser extension use cases	2718
Install and configure	2719
Organizational deployment	2720
Browser extension usage guidelines	2721
Browser extension known limitations	2722
Using Amazon Quick Flows in the browser extension	2722
Amazon Quick Microsoft Outlook extension	2724
Microsoft Outlook extension usage guidelines	2725
Sample prompts	2726
Amazon Quick Microsoft Excel extension	2726
Microsoft Excel extension usage guidelines	2727
Sample prompts	2728
Amazon Quick Microsoft Word extension	2728
Microsoft Word extension usage guidelines	2729
Sample prompts	2730
Amazon Quick Microsoft PowerPoint extension	2730
Microsoft PowerPoint extension usage guidelines	2731
Sample prompts	2732
Amazon Quick Slack extension	2732
Slack extension author guide	2733
Slack extension user guide	2737
Amazon Quick Microsoft Teams extension	2741
Microsoft Teams extension author guide	2742
Microsoft Teams extension user guide	2746
Amazon Quick on desktop	2750
What is Amazon Quick on desktop?	2751
What you can do with Quick on desktop	2751
Next steps	2752
Enterprise setup	2752

How enterprise sign-in works	2752
Supported token signing algorithms	2753
Prerequisites	2754
Setup process	2754
Managed deployment configuration	2755
Microsoft Entra ID	2759
Google Workspace	2766
Okta	2771
Ping Identity	2777
Troubleshooting	2784
MDM deployment	2787
Getting started	2796
Prerequisites	2797
Downloading and installing	2798
Signing in	2799
First-run setup	2800
Desktop interface	2800
Home	2800
Features	2800
Recents	2801
Account menu	2801
Overlay panels and top bar	2801
Working with chat	2802
Response preferences	2802
Chat input	2803
Citations and sources	2806
Managing messages	2806
Running work in the background	2806
Coding agent sessions	2807
Plan mode	2807
Deep research	2807
Incognito mode	2808
Managing conversation context	2808
Away From Keyboard (AFK) recap	2808
Managing conversations	2809
Creating documents and visuals	2809

Session tabs	2809
All data and apps	2810
Exporting a conversation	2810
Activity feed	2810
On Home	2810
Feed	2811
Briefing	2811
Filters and search	2812
Auto resolve	2812
Catch-up	2812
Always on	2812
Calendar	2812
To-dos	2813
Configuring the feed	2813
Library	2813
What Library holds	2813
Views and filters	2814
Item actions	2814
Sharing files	2814
Apps in Amazon Quick	2814
Spaces	2815
Browsing spaces	2815
Chatting with a space	2815
Creating and organizing spaces	2815
Customize	2815
Agents	2816
Creating an agent	2816
Agent configuration	2817
Using an agent	2817
Monitoring agents	2818
Managing agents	2819
Skills	2820
Skill definition	2820
How agents load skills	2821
Owner and viewer	2821
System skills	2821

Creating a skill	2821
Using a skill	2822
The skill detail panel	2822
Draft and published versions	2823
Managing skills	2824
Browsing for more skills	2824
Sharing your skills	2825
Administrator controls	2825
Connectors	2826
Discovering and adding connectors	2826
Managing your installed connectors	2827
Managing tool permissions	2829
MCP servers	2830
Coding agents	2832
Knowledge	2832
Indexing options	2833
Details	2833
Advanced settings	2833
Search indexing limits	2834
Folder access permissions	2834
Mission Control	2834
Why use Mission Control	2834
Capabilities	2834
Routines	2835
Working with agent runs	2837
KPIs and metrics	2837
Viewing your runs and creator analytics	2837
Accessing Mission Control	2838
Examples of routines	2838
My context	2839
Memories	2839
Reviewing and organizing memories	2839
Memory settings	2840
Knowledge graph	2840
Entity categories	2840
Data sources	2841

Viewing the knowledge graph	2841
Configuring and resetting the knowledge graph	2842
Settings	2842
Account	2843
Appearance and language	2843
Chat	2843
General	2843
Advanced	2844
System tools	2844
Managing permissions	2844
Available system tools	2845
Security and privacy	2846
How data is handled	2847
Data storage	2847
Folder permissions	2847
System tool permissions	2847
Connection security	2848
Privacy controls	2848
Network access and required domains	2848
Clearing all data	2849
Troubleshooting	2850
Installation and launch issues	2850
Connection issues	2850
Performance issues	2850
File access issues	2850
Browser automation issues	2851
Schedule issues	2851
Managing agent-hours consumption	2851
Session usage	2851
MCP server issues	2851
Exporting diagnostics	2852
Getting additional help	2852
Amazon Quick on mobile	2853
Getting the app	2853
Signing in	2853
Feature availability at a glance	2854

Chat	2854
Creating documents and visuals	2855
Your content and data	2856
Connectors, agents, and skills	2856
Settings	2857
Mobile-specific behaviors	2857
What is not available on mobile	2857
Differences between iOS and Android	2858
Document history	2859
Previous updates	2969
Attributions	2982

What is Amazon Quick?

Amazon Quick is an AI-powered service for automating tasks, analyzing data, building web applications, and conducting research. You interact with Quick through natural language chat, and Quick uses AI agents to process your requests against connected data sources and applications.

Quick includes the following features:

- **Amazon Quick Sight** – Interactive data visualization and business intelligence. Connect to data sources, build dashboards, and embed analytics in applications.
- **Amazon Quick Flows** – Automate repetitive tasks with AI-powered workflows that use your data and take actions in connected applications.
- **Amazon Quick Automate** – Build business process automations with AI agents that make contextual decisions and execute actions across applications.
- **Amazon Quick Index** – Connect your organization's documents and data sources so AI responses are grounded in your information.
- **Amazon Quick Research** – Conduct in-depth research across the web and your data, delivered as a cited report.
- **Apps in Amazon Quick** – Build interactive web applications by describing what you need in natural language.

Note

Amazon Quick evolved from Amazon QuickSight. QuickSight continues as Amazon Quick Sight, a feature within Quick. All existing QuickSight APIs, SDKs, and integrations continue to work without changes.

Benefits of Quick

Everything in one place

Agents, research, automation, analytics, and app building are all built in. You don't need separate products for each capability.

Fully managed

No infrastructure to provision, no models to host, no ML expertise required.

Connects to your tools and data

Quick integrates with the apps, databases, and systems your organization already uses through knowledge bases, action connectors, extensions, and structured data connections. Open standards including Model Context Protocol (MCP) and OpenAPI let you connect any system.

Meets you where you work

Use Quick from the web, the desktop application, or directly inside Chrome, Slack, Microsoft Teams, and Microsoft 365 applications.

AI you can trust

Quick runs on AWS with the same identity, network, and compliance controls your security team already approved. Your data stays in your control.

Accessing Quick

To start using Amazon Quick, you have two options:

1. **Sign up at aws.com/quick** – Use your email or social credentials (Google, Apple, Amazon, GitHub). No AWS account required. Free and Plus plans available.
2. **Provision through the AWS Management Console** – Configure Quick for your organization with IAM Identity Center or IAM Federation. Per-user subscriptions with standard AWS billing.

Getting started

To start using Quick immediately, sign up at aws.com/quick.

If your organization uses AWS:

1. [How Amazon Quick works](#) – Learn how Quick components work together.
2. [Setting up and signing into Amazon Quick](#) – Configure Quick for your organization.
3. [Getting started with Amazon Quick](#) – Start a conversation, explore features, and see what Quick can do.

How Amazon Quick works

When you use Quick, the following components work together to process your requests.

You start in chat

Chat is the primary interface for Quick. You type a question, give an instruction, or describe what you need. Quick interprets your request and determines how to respond: answering from connected data, generating content, running a workflow, conducting research, or taking action in an external application. For more information, see [Ask questions, explore data, and get insights with chat in Amazon Quick](#).

Agents process your requests

Behind every interaction is an agent. Agents are configured with instructions that define their behavior, knowledge sources that ground their responses, and tools that let them take action. Quick provides a default agent, and you can create custom agents for specific domains. For more information, see [Create, customize, and deploy AI-powered chat agents in Amazon Quick](#).

Spaces organize context

Spaces bring together the resources an agent needs: documents, dashboards, datasets, knowledge bases, and connectors. When you assign a space to an agent, it draws on everything in that space to answer questions and complete tasks. You share spaces with your team so everyone benefits from the same context. For more information, see [Organize, collaborate, and share resources with spaces in Amazon Quick](#).

Integrations connect Quick to your world

Integrations give Quick access to external information and the ability to act on your behalf:

Knowledge bases

Bring external content into Quick for AI retrieval. Sources include Amazon S3, SharePoint, OneDrive, Confluence, Google Drive, and web crawlers. Quick keeps the index in sync as sources update.

Connectors

Let Quick read data, trigger workflows, and update records in external services. You create connectors from OpenAPI specifications or Model Context Protocol (MCP) servers.

Extensions

Embed Quick inside Chrome, Slack, Microsoft Teams, and Microsoft 365 applications so you interact with agents without leaving the app you're working in.

Structured data connections

Connect Quick Sight to databases, data warehouses, and data lakes for analytics and visualization.

For more information, see [Work with integrations in Amazon Quick](#).

Features use those building blocks

Each Quick feature combines the components above in different ways:

Feature	What it does	What it uses
Quick Sight	Interactive dashboards and analytics	Structured data connections, SPICE, datasets
Quick Flows	Task automation	Connectors, spaces, agent logic
Quick Automate	End-to-end process automation	Connectors, agents, human-in-the-loop
Quick Research	In-depth cited reports	Spaces, knowledge bases, web
Apps in Quick	Interactive web applications	Structured data, connectors, Quick Sight visuals
Desktop application	Personalized AI on your machine	Local files, email, calendar, MCP servers

Where to go next

I want to...	Go to
Set up Quick for my organization	Setting up and signing into Amazon Quick
Start using Quick	Getting started with Amazon Quick
Build dashboards and visualize data	Visualize, analyze, and share data with analyses, dashboards, and reports in Amazon Quick Sight
Automate tasks	Using Amazon Quick automations
Research a topic	Using Amazon Quick Research
Build an app	Build web applications with apps in Amazon Quick
Connect Quick to my tools	Work with integrations in Amazon Quick
Use Quick on my desktop	What is Amazon Quick on desktop?

Setting up and signing into Amazon Quick

Amazon Quick offers two ways to get started, depending on your needs:

- **Get started quickly** – If you're an individual or small team and don't have an AWS account, see [Signing up at aws.com/quick](https://aws.com/quick).
- **Use Quick with your organization's AWS account** – If your organization already uses AWS services, see [Signing up through the AWS Console](#).

Free and Plus accounts: Sign up at <https://aws.com/quick> using an email address. Amazon Quick creates an account for you automatically using your email or social login credentials. Choose from Free, Free Trial Plus, or Paid Plus plan tiers. No AWS account or technical configuration is required.

AWS Console accounts: If your organization uses AWS, you can provision Amazon Quick through the AWS Management Console. This path uses IAM Identity Center for authentication and integrates with your existing AWS billing and governance.

If your IAM Identity Center instance is replicated to multiple AWS Regions, you can create your Quick Enterprise subscription in an active additional Region. For more information, see [Set up Amazon Quick with IAM Identity Center multi-Region](#).

For a comparison of features available with each account type, see [Pricing and availability](#).

For more information on administering your account after setup, see [Administering Amazon Quick \(Free/Plus/Max\)](#) or [Administering Amazon Quick](#) for accounts provisioned through the AWS Management Console.

Topics

- [Complete initial configuration tasks](#)
- [Signing up at aws.com/quick](https://aws.com/quick)
- [Signing up through the AWS Console](#)
- [Signing in to Amazon Quick](#)

Complete initial configuration tasks

Note

The following configuration tasks apply to AWS Console accounts only. If you are signing up for Free/Plus/Max with email at aws.com/quick, these steps are handled automatically. See [Signing up at aws.com/quick](https://aws.com/quick) instead.

To use Amazon Quick you must first complete the following tasks:

Topics

- [Sign up for an AWS account](#)

Sign up for an AWS account

To get started with AWS, you need an AWS account. For information about creating an AWS account, see [Getting started with an AWS account](#) in the *AWS Account Management Reference Guide*.

Signing up at aws.com/quick

Amazon Quick is available directly at <https://aws.com/quick> – no existing AWS account required. This page walks you through creating a Amazon Quick Free or Plus account using your email address or social login credentials.

Note

If your organization already uses AWS and you want to access Amazon Quick through the AWS Management Console, see [Signing up through the AWS Console](#) instead. For a comparison of features available with each account type, see [Pricing and availability](#).

Topics

- [Who is this for?](#)
- [Sign-in options](#)
- [Step-by-step signup procedure](#)

- [What happens behind the scenes](#)

Who is this for?

The direct signup at aws.com/quick is designed for:

- **Individuals** who want to explore Amazon Quick without creating an AWS account.
- **Small teams** who need a fast way to get started with AI-powered productivity tools.
- **Anyone** who wants to try Amazon Quick with a few clicks.

You can sign up with an email address or use an existing Google, Apple, GitHub, or Amazon account. An AWS account is created and managed automatically behind the scenes – you never need to interact with it directly.

Sign-in options

When you sign up at aws.com/quick, you can create your account using:

- **Email address** – Create an account with any email (covered in this guide).
- **Google** – Sign in with your Google account.
- **Apple** – Sign in with your Apple ID.
- **GitHub** – Sign in with your GitHub account.
- **Amazon** – Sign in with your Amazon account.

This guide covers the email-based signup flow. Social login options follow a similar process but use your existing credentials from the selected provider.

Step-by-step signup procedure

To sign up for a Amazon Quick Free or Plus account

1. **Navigate to aws.com/quick.**

Open your web browser and go to <https://aws.com/quick>. You see the Amazon Quick landing page with an option to **Sign up for free**.

2. **Choose Sign up for free.**

Choose **Sign up for free** to begin the account creation process. You are redirected to the AWS signup page.

3. Enter your email address.

In the email field, type the email address you want to use for your Amazon Quick account. This email becomes your sign-in identity for Amazon Quick.

After entering your email, choose **Continue**.

Tip

Use an email address you check regularly. You need to verify it in the next step, and Amazon Quick sends important account notifications to this address.

4. Enter your name.

On the next screen, enter your name. This is the display name that other users in your Amazon Quick account see.

Your email address is shown for confirmation. If you need to change it, choose **Change** to go back.

Choose **Continue** to proceed.

5. Verify your email address.

Amazon Quick sends a 6-digit verification code to the email address you provided. Check your inbox (and spam/junk folder if needed) for an email from AWS.

Enter the 6-digit code in the verification field and choose **Continue**.

Note

If you don't receive the code, choose **Resend code** to send a new one. The resend option becomes available after a brief countdown.

6. Create a password.

Create a password for your account. Your password must meet standard AWS security requirements:

- At least 8 characters
- At least one uppercase letter
- At least one lowercase letter
- At least one number
- At least one special character

Enter your password, then re-enter it to confirm. Choose **Continue**.

7. Wait for account provisioning.

After setting your password, Amazon Quick provisions your account. You see a "Setting up" screen with a brief message. This process typically takes less than one minute.

Note

During this step, Amazon Quick automatically creates and configures an AWS account on your behalf. You don't need to manage this account – it's handled entirely by the service.

8. Choose how to access Amazon Quick.

Once provisioning completes, the Welcome page appears. You have two options:

- **Download for macOS** – Download the Amazon Quick desktop application for your computer.
- **Continue on web for now** – Start using Amazon Quick immediately in your web browser.

Choose the option that works best for you. You can always download the desktop app later from the **Quick anywhere** menu in your profile settings.

9. Choose your work focus (Getting Started wizard).

Amazon Quick presents a Getting Started wizard that helps personalize your experience. You can select one or more work focus topics, or skip this step entirely.

The available work focus topics are:

Topic	Description
Communicate with stakeholders	Tools and templates for stakeholder communication
Grow my sales pipeline	Sales pipeline management and analysis
Improve my business performance	Business performance analytics and insights
Organize my team's knowledge	Knowledge management and documentation
Plan and prioritize my work	Project planning and task prioritization
Build and ship my product	Product development and delivery workflows
Improve my marketing reach	Marketing analytics and content creation
Hire and develop my team	HR, recruiting, and team development
Manage risk and compliance	Risk assessment, compliance, and governance

Select a topic that matches your primary use case, then choose **Let's go**. Alternatively, choose **Skip and start chatting** to go directly to the chat interface.

10. Start using Amazon Quick.

The Amazon Quick chat interface loads. If you selected a work focus topic, you see personalized suggestion cards tailored to your selection. The interface includes:

- **Left navigation** – Access New chat, Search, Spaces, Research, Chat agents, and More.
- **Chat input** – Type questions or requests in the **Ask me anything** field.
- **Attach files** – Upload documents for analysis (up to 20 files, under 50 MB each).
- **Suggestion cards** – Pre-built prompts based on your selected work focus.

You're now ready to start using Amazon Quick. For information about features and capabilities, see [Getting started with Amazon Quick](#).

What happens behind the scenes

When you sign up at aws.com/quick, Amazon Quick automatically:

1. Creates a **sign-in identity** linked to your email address – this is what you use to sign in to Amazon Quick.
2. Provisions an **AWS account** that hosts your Quick resources – you don't need to access or manage this account.
3. Sets up your Quick environment with default settings and built-in AI agents.

Signing up through the AWS Console

Note

This section covers signing up for Amazon Quick through the AWS Management Console for organizations that already use AWS. If you want to sign up without an AWS account, see [Signing up at aws.com/quick](https://aws.com/quick) for a faster setup process using email or social login. For a comparison of features available with each account type, see [Pricing and availability](#).

When you first sign up for Amazon Quick, you get a free trial subscription for twenty-five users for 30 days. During the process of signing up, you may set options for your identity provider.

Before you begin, make sure that you can connect to an existing AWS account. If you don't have an AWS account, see [Complete initial configuration tasks](#). The person who signs up for Quick needs to have the correct AWS Identity and Access Management (IAM) permissions. For more information, see [IAM policy examples for Quick](#).

To test your permissions, you can use the IAM policy simulator; for more information, see [Testing IAM policies with the IAM policy simulator](#). Also, check whether your AWS account is part of an organization based on the AWS Organizations service. If so and you sign in as an IAM user, make sure that you didn't inherit any IAM permissions that deny access to the required permissions. For more information on Organizations, see [What is AWS Organizations?](#)

 Note

Your data is encrypted by default using AWS-managed keys. Admins can adjust settings for custom encryption in the admin portal after signing up.

To subscribe to Quick

1. Sign in to your AWS account and open Quick from the AWS Management Console. You can find it under **Analytics** or by searching for *Quick*.

Your AWS account number is displayed for verification purposes.

2. Enter a unique account name for Quick.
 - Enter a notification email address for the Quick account owner or group. This email address receives service and usage notifications.
3. Choose the AWS Region that you want to use for your initial data storage capacity, called SPICE.
4. Choose an authentication method that you want to connect to Quick with. Select from one of the following:
 - **(Recommended) Password-based or Single-Sign On**
 - **IAM Identity Center**
 - **Single-Sign On Only**
 - **Active Directory**
5. Review the choices that you made, then choose **Create account**.
6. Upon completion, your Quick account is created. To open Quick, choose **Go to Quick**.

Non-admin sign-up process

If you are not an administrator, you can sign up for Quick using a link provided by your admin. This process allows you to create your account and configure your Quick environment.

To sign up as a non-admin user

1. **Getting your sign-up link**

Sign up for Amazon Quick with the link provided by your admin.

2. Account creation and Quick access

After you've created your account, from within Quick, search for Amazon Quick.

3. Account configuration

Enter your Account name, Email, Region, and Authentication method.

- For authentication, Amazon Quick supports Pro license integration with start URL format:
`https://amzn.awsapps.com/start`
- Default Region is us-east-1

4. Plan selection

Select the plan and storage options you are using for your instance of Amazon Quick.

Signing in to Amazon Quick

You can sign in to Amazon Quick multiple ways, depending on what your Quick administrator has set up. You can sign in to Quick using AWS root, AWS Identity and Access Management (IAM), corporate Active Directory, or your native Quick credentials. If your Quick account is integrated with an identity provider such as Okta, the following procedures don't apply to you.

If you're a Quick administrator, make sure to allow-list the following domains within your organization's network.

User type	Domain or domains to allow-list
Users who sign in directly through Quick and Active Directory users	<code>signin.aws</code> and <code>awsapps.com</code>
AWS root user	<code>signin.aws.amazon.com</code> and <code>amazon.com</code>
IAM users	<code>signin.aws.amazon.com</code>

Signing in to a Amazon Quick Free or Plus account

If you created your Amazon Quick account at aws.com/quick, follow these steps to sign in.

To sign in to a Amazon Quick Free or Plus account

1. Navigate to the [sign-in page](#).
2. Choose **Sign in**.
3. Enter the email address associated with your account.
4. Enter your password.
5. Choose **Sign in**.

If you have multiple Amazon Quick accounts associated with your email, you may be prompted to select which account to access. Choose the account you want to sign in to.

Tip

You can also sign in using the social login option you used during signup (Google, Apple, GitHub, or Amazon) if you linked your account during registration.

Note

If you signed up through the AWS Management Console instead, see the next section for console sign-in instructions.

Post-setup configuration

After signing in to Quick, you can customize your environment by configuring agent settings and other preferences.

To complete post-setup configuration

1. **Agent customization access**

From the left navigation menu, under Customization, select Agent customization.

2. **Environment settings**

Enter the settings you require for your environment. These settings will customize Amazon Q's behavior for your specific use case.

Important

We strongly recommend that you don't use the AWS root user for your everyday tasks, even the administrative ones. Instead, adhere to the best practice of using the root user only to create your first IAM user. Then securely lock away the root user credentials and use them to perform only a few account and service management tasks. For more information, see [AWS account root user](#) in the *IAM User Guide*.

Signing in through the AWS Console

Use the following procedure to sign in to Quick through the AWS Management Console.

To sign in to Quick

1. Go to <https://quicksight.aws.amazon.com/>.
2. For **Quick account name**, enter your account name. This is the name that was created when the Quick account was created in AWS.

If you were invited to the Quick account by email, you can find the account name inside of that email. If you don't have the email that invited you to Quick, ask the Quick administrator in your organization for the information that you need.

You can also find your Quick account name by choosing the profile icon on the Quick console menu. In some cases, you might not have access to your Quick account or have an administrator who can provide this information, or both. If so, contact AWS Support and open a ticket that includes your AWS customer ID.

3. For **Username**, enter your Quick user name. User names that contain a semicolon (;) aren't supported. Choose one of the following:
 - For organizational users – The user name is provided by your administrator.

Your account can be based on IAM credentials or your email address if it's a root email address. Or it can be used as the user name to invite you into the Quick account. If you

received an invitation email from another Quick user, it indicates what type of credentials to use.

- For individual users – The user name that you created for yourself.

This is usually the IAM credentials that you created.

The remaining steps vary depending on the user type you sign in as (directly through Quick or as an Active Directory user, AWS root user, or IAM user). For more information, see the following sections.

Finishing Quick sign-in as a Quick or Active Directory user

If you're signing in directly through Quick or are using your corporate Active Directory credentials, you're redirected to `signin.aws` after you enter your account name and user name. Use the following procedure to finish signing in.

To finish signing in to Quick if you sign in directly through Quick or use Active Directory credentials

1. For **Password**, enter your password.

Passwords are case-sensitive and must be 8–64 characters in length. They must also contain each of the following:

- Lowercase letters (a–z)
- Uppercase letters (A–Z)
- Numbers (0–9)
- Nonalphanumeric characters (~!@#\$\$%^&* _-+= `| \(){}[];:"'<>.,?/)

2. If your account is multi-factor authentication enabled, enter the multi-factor authentication code that you receive for **MFA code**.
3. Choose **Sign in**.

Finishing Quick sign-in as an AWS root user

If you're signing in as an AWS root user, you're redirected to `signin.aws.amazon.com` (or `amazon.com`) to complete the sign-in process. Your user name is prefilled. Use the following procedure to finish signing in.

To finish signing in as an AWS root user

1. Choose **Next**.
2. For **password**, enter your password. For more information about root user passwords, see [Changing the AWS account root user password](#) in the *IAM User Guide*.
3. Choose **Sign in**.

Finishing Quick sign-in as an IAM user

If you're signing in as an IAM user, you're redirected to signin.aws.amazon.com (or amazon.com) to complete the sign-in process. Your user name is prefilled. Use the following procedure to finish signing in.

To finish signing in as an IAM user

1. For **Password**, enter your password. For more information about IAM user passwords, see [Default password policy](#) in the *IAM User Guide*.
2. Choose **Sign in**.

Signing in with multiple accounts

You can sign in to Amazon Quick with multiple accounts at the same time and switch between them as necessary.

To sign in to multiple accounts

1. Follow the steps at [Signing in to Amazon Quick](#) to sign in to Amazon Quick.
2. From the navigation bar, choose your account name to reveal a dropdown menu.
3. Hover over **Switch accounts**. A list of accounts that you've signed in to before appears. Do one of the following:
 - To switch to an account, select that account and provide sign-in credentials.
 - To sign in to a new account, choose **Add another account**. You are prompted to provide sign-in credentials for the account.

To sign out of an account or all accounts

1. From the navigation bar, choose your account name to reveal a dropdown menu.
2. Hover over **Sign out**. You can choose to sign out of your current account or to **Sign out of all accounts**.

URL structure

When using multiple accounts, Amazon Quick URLs include the account name to identify which account you're accessing.

Application URLs

Application URLs follow this structure:

```
https://us-east-1.quicksight.aws.amazon.com/sn/account/ACCOUNT-NAME/...
```

Examples:

```
https://us-east-1.quicksight.aws.amazon.com/sn/account/my-account/start/home  
https://us-east-1.quicksight.aws.amazon.com/sn/account/my-account/dashboards/1234
```

Console URLs

Console URLs follow this structure.

Standard accounts (account alias is the account ID)

```
https://us-east-1.quicksight.aws.amazon.com/sn/console/account/123456789012/...
```

Examples:

```
https://us-east-1.quicksight.aws.amazon.com/sn/console/account/123456789012/admin  
https://us-east-1.quicksight.aws.amazon.com/sn/console/account/123456789012/asset-  
management
```

Accounts with IAM multi-account access (account alias is the session differentiator)

```
https://us-east-1.quicksight.aws.amazon.com/sn/console/account/123456789012-abc123/...
```

Examples:

```
https://us-east-1.quicksight.aws.amazon.com/sn/console/account/123456789012-abc123/admin
https://us-east-1.quicksight.aws.amazon.com/sn/console/account/123456789012-abc123/asset-management
```

Key considerations and limitations

- The URL structures above enable bookmarking of resources within the Amazon Quick account.
- You can simultaneously be signed in to different accounts in different tabs or windows in your browser.
- The Amazon Quick URL contains your account name. If you change it to another account name, you are prompted to sign in if you aren't already.
- You can sign in to up to five accounts. If you try to sign in to a sixth account, the system automatically signs out and forgets the least recently signed in account.
- Multi-account sign-in is not supported for IAM users who access Amazon Quick through IAM-based multi-account federation.

IAM SSO federation with multi-account

For more information about setting up identity federation, see [Initiating sign-on from Quick](#).

Service Provider (SP) initiated federation

If your account has SSO configuration established in Quick, all login attempts or links that include the account alias provide seamless login. There is no change in behavior for these logins.

Identity Provider (IdP) initiated federation

Scenarios with no change in behavior:

- If you have enabled IAM multi-account access, all federation links provide seamless login, regardless of whether they include the account alias.
- If you have SSO configuration enabled in Quick, all IdP initiated federation or IdP deep links that include the account alias provide seamless login.
- If you have never used Quick before or have only ever used one account with Quick and you use IdP federation to access that one account, there is no change in experience.

Scenarios with change in behavior: When you have one or more saved Quick accounts and access any IdP federation link without an account alias, instead of being taken directly to your account, you land on the Quick saved account login page where you can select your session and be taken to your account.

Recommended configuration: To ensure seamless IdP federation, include the account alias in all IdP federation links. If you are unable to ensure the account alias in your deep links or you do not have SSO configuration enabled in Quick, include the `sso_login=true` query parameter in your links.

Recommended example start federation URL:

```
https://us-east-1.quicksight.aws.amazon.com/sn/account/ACCOUNT-NAME/start/home?  
sso_login=true
```

If your sign-in process happens automatically and you need to use a different account, use a private or incognito browser window. Doing this prevents the browser from reusing cached settings.

 Note

If you're not sure which type of account you have: if you signed up at aws.com/quick and use an email/password to sign in, you have a Free or Plus account. If you sign in through the AWS Management Console or your organization's single sign-on portal, you have an AWS Console account.

Getting started with Amazon Quick

Amazon Quick is an AI-powered workspace that combines business intelligence, AI-assisted chat, workflow automation, and research capabilities in a single service. Use the topics in this section to get started with the core capabilities of Quick, including AI agents for natural language conversations, Amazon Quick Sight for data visualization, Spaces for team collaboration, Quick Flows for workflow automation, and Research for comprehensive analysis.

You can begin by exploring the platform's features through AI chat, using sample data for visualizations, or connecting to your own data sources to start building real-world solutions.

Topics

- [Using the Amazon Quick console](#)
- [Getting started with Amazon Quick Sight](#)

Using the Amazon Quick console

In the following topic, you can find a brief introduction to using the Amazon Quick user interface.

Topics

- [Using the Amazon Quick menu and landing page](#)
- [Searching Amazon Quick](#)
- [Choosing a language in Amazon Quick](#)

Using the Amazon Quick menu and landing page

After you sign in to Amazon Quick, you see the Amazon Quick landing page. This page provides tabs for your analyses, dashboards, folders, data, and more. It also provides a menu bar at the top, with options for the following:

- Searching Amazon Quick
- Accessing your user profile (managing Quick, language selection, and help)
- Interacting with the Quick chat agent
- Choosing the AWS Region that you want to work in

- Discovering helpful resources (documentation, how-to videos, community forum, blogs, and more)

Note

Consult your administrator before changing your AWS Region. Your default AWS Region is configured by your Amazon Quick administrator. Changing the AWS Region changes where your work is stored.

Navigating the homepage

Using the navigation menu on the left, you can explore how to work with dashboards, build powerful analyses, manage your data sources, organize content with folders, leverage advanced analytics capabilities, and more.

The following are all the available Amazon Quick capabilities:

- **Agents** – Build AI agents that intelligently analyze, plan, and work with your data to help complete tasks.
- **Spaces** – Work with Amazon Q using specific files, data, and information to complete projects or organize teams.
- **Scenarios** – Elevate your data exploration in the AI-driven workbook canvas that is enhanced with generative analysis capabilities. Allow AI to suggest relevant data to enrich your insight discovery.
- **Stories** – Use Amazon Q prompts and visuals to produce a story that incorporates the details that you provide.
- **Topics** – Create a new generative topic to introduce the generative Q&A experience for users.
- **Dashboards** – Share interactive data visualizations with your audience. You can explore and filter visualizations according to your specific needs and generate snapshots for easy sharing.
- **Analyses** – Dive deep into your data with robust visualization and analysis tools. Create charts and graphs to uncover insights before publishing to a wider audience.
- **Favorites** – Favorite the dashboards and analyses you want to easily find later.
- **Folders** – Organize your analyses and dashboards in folders.
- **Automations** – Create automated processes that reduce manual intervention and save valuable time.

- **Flows** – Create and share intelligent workflows to accomplish goals quickly and predictably.
- **Data** – Connect to data and visualize, create specialized chat assistants, or mix both.
- **Actions** – Perform tasks such as creating and updating tickets, managing incidents, and accessing project information directly.

Customizing the homepage layout

The Quick homepage provides a customizable workspace where you can add widgets that enhance your analytics experience. The homepage features three key widget types:

- **Quick use cases** – Offers quick access to tutorial videos and onboarding resources.
- **Recents** – Displays recently viewed resources for faster navigation.
- **Business Metrics** – Allows you to pin important visualizations and metrics directly from your dashboards.

You can easily add, remove, and rearrange these widgets to create a personalized homepage that suits your workflow.

To add a widget

1. From the Quick homepage, choose **Add widget**.
2. You can choose to add either **Quick use cases**, **Recents**, or **Business metrics**.
 - a. If **Quick use cases** is already added to the homepage, this option will be unavailable. If you previously deleted the widget, choose **Add widget** then **Quick use cases** to add it back to the homepage. **Quick use cases** always appears at the top of the homepage to ensure easy access to learning resources, even if previously dismissed.
 - b. If you choose **Recents**, a widget containing recently viewed resources is added to the homepage. Choose a resource to quickly pick up where you left off.
 - c. If you choose **Business metrics**, you will be redirected to the **Dashboards** page. From there, you can open a dashboard and choose which visual you would like to pin to the homepage by selecting the home + icon.

To delete a widget

1. Navigate to the widget you would like to delete.

2. Choose the three-dot menu on the top right of the widget.
3. Choose **Delete**.
4. The selected widget should be removed from the homepage.

To reorder a widget

1. Select the widget you would like to reorder.
2. When the widget is selected, the handle icon appears on the top left of the widget.
3. Choose the handle icon and drag the widget to swap its place with another widget.

To resize a widget

1. Select the widget you would like to resize.
2. When the widget is selected, a resize icon appears on the bottom right.
3. Choose the resize icon and drag the widget up or down to reach your desired size.

Accessing user profiles

To access the user profile menu, choose the user icon at the upper right of any page in Amazon Quick. Use this menu to manage Amazon Quick features, choose a language, change your AWS Region, get help from the documentation, or sign out of Amazon Quick.

The following options are available from the user profile menu:

- **Manage Quick** – If you have appropriate permissions, you can access administrative functions such as managing users, subscriptions, [SPICE](#) capacity, and account settings.
- **Language setting** – Choose the language you want to use in the Amazon Quick user interface.
- **Region setting** – Choose the AWS Region that you want to work in.

Note

Consult your administrator before changing your AWS Region. Your default AWS Region is configured by your Amazon Quick administrator. Changing the AWS Region changes where your work is stored.

- **Help** – This will open the official Quick documentation, which you can view online, in Kindle, or as a PDF.
- **Sign out** – Choose this option to sign out of Amazon Quick and your AWS session.

Searching Amazon Quick

From the search bar, you can search for analyses and dashboards. To use the search tool, go to the **Start Page** and choose the search box at the top-left of the page. Then enter the name, or part of the name, of the data set, analyses, or dashboard you want to find. The search is not case-sensitive.

After you locate the item that you're looking for, you can open it directly from the search results. You can modify a dataset, create an analysis from a dataset, or access an analysis or dashboard. Choose an item from the search results to open it.

Choosing a language in Amazon Quick

You can choose the language that you want to use in the Amazon Quick user interface. This option is set separately for each individual user. The first time a user signs in, Amazon Quick detects and selects a suitable language. This choice is based on the user's browser preferences and interactions with localized AWS websites.

Quick supports the following languages:

Languages available in the Amazon Quick user interface

Official name	Language code	Localized name
Dansk	da-DK	Danish
Deutsch (Deutschland)	de-DE	German
English (US)	en-US	English
Español (España)	es-ES	Spanish
Français (France)	fr-FR	French

Official name	Language code	Localized name
Italiano (Italia)	it-IT	Italian
Nederlands (Nederland)	nl-NL	Dutch
Norsk bokmål	nb-NO	Norwegian
Português (Brasil)	pt-BR	Portuguese
Suomi	fi-FI	Finnish
Svenska (Sverige)	sv-SE	Swedish
日本語	ja-JP	Japanese
한국어	ko-KR	Korean
中文 (中国)	zh-CN	Simplified Chinese
中文 (台灣)	zh-TW	Traditional Chinese

Choosing a language translates only user interface elements. It doesn't translate the following:

- Amazon Quick reserved keywords
- User input
- Data
- Date or number formats
- ML Insights, suggested insights, or computations in narratives (including text)

Use the following procedure to change the language in the Amazon Quick interface.

1. Choose your account name at the top right.

2. Choose **Preferences**.
3. Choose the language that you want to use from the **Preferred Language** dropdown list.

Getting started with Amazon Quick Sight

Amazon Quick Sight is the data visualization and analysis component of Amazon Quick. This section walks you through creating your first analysis using sample data, including uploading datasets, creating visualizations with AutoGraph, and sharing your results through dashboards.

Topics

- [Quick start: Create an Amazon Quick Sight analysis with a single visual using sample data](#)
- [Tutorial: Create an Amazon Quick Sight dashboard using sample data](#)
- [Sample data](#)

Quick start: Create an Amazon Quick Sight analysis with a single visual using sample data

Before you create your first analysis, make sure to complete the steps in [Setting up and signing into Amazon Quick](#).

With the following procedure, you use the Web and Social Media Analytics sample dataset to create an analysis containing a line chart visual. This visual shows the count by month of people that have added themselves to the mailing list.

To create an analysis containing a line chart visual using a sample dataset

1. From the Amazon Quick homepage, from Amazon Quick Sight, choose **Analyses** from the left navigation menu. If you don't have sample data, you can download it from [web-and-social-analytics.csv.zip](#). Unzip the file so you can use the .csv file.

To upload the sample data, do the following:

- a. Choose **Data** from the left navigation menu. Under the **Dataset** tab, select **New** then **Dataset**.
- b. Choose **Upload file**.

- c. Choose the sample file, `web-and-social-analytics.csv`, from your drive. If you don't see it, check that you unzipped the `web-and-social-analytics.csv.zip` file.
 - d. Confirm file upload settings by choosing **Next** on the **Confirm file upload settings** screen.
 - e. Choose **Visualize** on the **Data source details** screen.
 - f. Skip the next step. Choosing **Visualize** brings you to the same screen as the process in Step 2.
2. On the **Datasets** page, choose the **Web and Social Media Analytics** dataset, and then choose **Use in Analysis**.
3. In the **Data** pane, choose **Date**, and then choose **Mailing list adds**.

Amazon Quick Sight uses AutoGraph to create the visual, selecting the visual type that it determines is most compatible with those fields. In this case, it selects a line chart that shows mailing list adds by day, which is the date granularity default.

4. Navigate to the **Field wells** at the bottom of the **Visuals** pane.
5. Choose the **X axis** field well. Select the three-dot menu, choose **Aggregate**, and then choose **Month**.

The line chart updates to show mailing list adds by month, rather than by the default of by year.

Tutorial: Create an Amazon Quick Sight dashboard using sample data

Use the procedures in the following sections to complete these tasks:

- Create and prepare a Marketing dataset using the Web and Social Media Analytics sample data.
- Create a Marketing analysis and add several visuals to it.
- Modify the visuals in the analysis, including the following:
 - Adding another measure to an existing visual
 - Changing chart colors
 - Changing date granularity
 - Changing the size and layout of the visuals
 - Applying a filter
- Publish a dashboard based on the analysis.

Topics

- [Tutorial: Create a prepared Amazon Quick Sight dataset](#)
- [Tutorial: Create an Amazon Quick Sight analysis](#)
- [Tutorial: Modify Amazon Quick Sight visuals](#)
- [Tutorial: Create an Amazon Quick Sight dashboard](#)

Tutorial: Create a prepared Amazon Quick Sight dataset

Use the following procedure to prepare the Marketing dataset and create an analysis. If you don't see the Web and Social Media Analytics sample data already in Amazon Quick Sight, you can download it: [web-and-social-analytics.csv.zip](#).

To prepare the Marketing dataset and create an analysis

1. From the Amazon Quick homepage, choose **Data** in the navigation pane.
2. In the **Datasets** tab, choose **New** then **Dataset**.
3. From the existing data sources, choose the **Web and Social Media Analytics** Amazon S3 data source. Choose your table and then choose **Edit/Preview data**.

Amazon Quick opens the data preparation page.

4. For **Dataset Name**, enter **Marketing Sample** to replace *Web and Social Media Analytics* for the dataset name.
5. Exclude some fields from the dataset.

In the **Fields** pane, choose the three-dot menu for the **Twitter followers cumulative** and **Mailing list cumulative** fields, and then choose **Exclude field**. To select more than one field at a time, press the Ctrl key while you select (Command key on Mac).

6. Rename a field.

In the **Dataset** preview pane, scroll to the **Website Pageviews** field and choose the pencil icon to edit.

In the **Edit field** page that opens, for **Name**, enter **Website page views**, and then choose **Apply**.

7. Add a calculated field that substitutes a text string for any 0-length string value in the **Events** field:

- a. On the data preparation page, scroll to the top of the **Fields** pane, and then choose **Add calculated field**.
- b. In the **Add calculated field** page that opens, for **Add name**, enter **populated_event**.
- c. In the **Functions** pane at right, double-click the **ifelse** function from the list of functions. This adds the function to the calculated field formula.
- d. Expand the **Field list** pane by choosing the drop-down arrow, and then double-click the **Events** field. This adds the field to the calculated field formula.
- e. In formula editor, enter the following additional functions and parameters required, in bold in the following: `ifelse(strlen({Events})=0, 'Unknown', {Events})`.

The final formula should be as follows: `ifelse(strlen({Events})=0, 'Unknown', {Events})`.

- f. Choose **Save**.

The new calculated field is created, and appears at the top of the **Fields** pane.

8. Choose **Save**.

Next steps

Create an analysis by using the procedure in [Tutorial: Create an Amazon Quick Sight analysis](#).

Tutorial: Create an Amazon Quick Sight analysis

In the following short tutorial, you create an analysis, add a visual using AutoGraph, and add another visual by choosing a specific visual type. This procedure builds on the dataset that you create and prepare in [Tutorial: Create a prepared Amazon Quick Sight dataset](#).

Create your analysis

Use the following procedure to create your analysis.

To create your analysis

1. On the Amazon Quick Sight homepage, choose **Analyses**.
2. Choose **Create analysis** and select the dataset to use.

Create a visual by using AutoGraph

Create a visual by using AutoGraph, which is selected by default.

On the analysis page, choose **Date** and **Return visitors** in the **Fields list** pane.

Amazon Quick Sight creates a line chart using this data.

Create a scatter plot visual

Create a visual by choosing a visual type and dragging fields to the field wells.

To create a scatter plot visual

1. On the analysis page, choose **Insert** and then **Add visual** on the application bar. A new, blank visual is created, and AutoGraph is selected by default.
2. In the **Visual types** pane, choose the scatter plot icon.
3. Choose fields in the **Fields list** pane to add to the **Field wells** pane:
 - Choose **Desktop Uniques** to populate the **X axis** field well.
 - Choose **Mobile Uniques** to populate the **Y axis** field well.
 - Choose **Date** to populate the **Group/Color** field well.

A scatter plot is created using these fields.

Next steps

Modify the visuals in your analysis by using the procedure in [Tutorial: Modify Amazon Quick Sight visuals](#).

Tutorial: Modify Amazon Quick Sight visuals

Use the following procedures to modify the visuals that you created using the procedures in [Tutorial: Create an Amazon Quick Sight analysis](#).

Modify the line chart visual

Modify your line chart visual by making it show an additional measure by date, and also by changing the chart color.

To modify your line chart visual

1. In your analysis, select the line chart visual.
2. Add another measure to the visual.

Select the **New visitors SEO** field in the **Fields list** pane. This measure is added to the **Value** field well, and the line chart updates with a line to represent it. The visual title also updates.

3. Change the color of the line used to represent the **Return visitors** measure.

Choose the line on the chart that represents **Return visitors**. To do this, choose the end of the line, not the middle of the line.

Choose **Color Return visitors**, and then choose the red icon from the color selector.

4. Choose the **Date** field in the **X axis** field well, choose **Aggregate**, and then choose **Month**.

Modify the scatter plot visual

Modify your scatter plot visual by changing the data granularity.

To modify your scatter chart visual

1. In the analysis, select the scatter plot visual.
2. Choose the **Group/Color** field well, choose **Aggregate**, and then choose **Month**.

The scatter plot updates to show the measures by month, rather than by the default of by year.

Modify both visuals by changing visual layout and adding a filter

Modify both visuals by changing visual size and location, and by adding a filter and applying it to both of them.

Change the visual layout

Modify both visuals by changing visual size and location.

To modify both visuals

1. In your analysis, select the line chart visual.

2. Choose the resize handle in the lower right corner of the visual and drag up and to the left, until the visual is half its former size both horizontally and vertically.
3. Repeat this procedure on the scatter plot visual.
4. Choose the move handle on the scatter plot visual, and drag it up to the right of the line chart visual so that they are side by side.

Modify both visuals by adding a filter

Modify both visuals by adding a filter and applying it to both of them.

To add a filter to both visuals

1. In the analysis, choose the scatter plot visual.
2. Choose **Insert** then **Add Filter** on the application bar.
3. Choose the **Date** field to filter on.
4. Select the new filter to expand it.
5. In the **Edit filter** pane, for **Condition**, choose the **After** comparison type.
6. Enter a start date value of 1/1/2014.

Choose **Date**, choose **2014** for the year, **January** for the month, and then choose **1** on the calendar.

7. In the **Edit filter** pane, choose **Apply** to apply the filter to the visual.

The filter is applied to the scatter plot visual. This is indicated with a filter icon on the visual drop-down menu.

8. Apply the filter to the line chart visual.

In the **Filter** pane, choose the **Date** filter again and choose **Single visual**, and then choose **All visuals of this dataset**.

The filter is applied to the line chart visual as well.

Next steps

Create a dashboard from your analysis by using the procedure in [Tutorial: Create an Amazon Quick Sight dashboard](#).

Tutorial: Create an Amazon Quick Sight dashboard

Use the following procedure to create a dashboard from the analysis that you created using the procedure in [Tutorial: Create an Amazon Quick Sight analysis](#).

To create a dashboard from your analysis

1. In your analysis, choose **Publish**, then choose **Publish dashboard**.
2. Choose **Publish as a new dashboard** or **Replace existing dashboard**.
3. For a new dashboard, enter **Marketing Dashboard** in **Dashboard name**. Optionally, add notes.
4. Under **Select sheets**, choose **All sheets** or **Specific sheets**.
5. Configure the following options as needed:
 - **Dashboard options - Enable Quick Actions** to allow dashboard users to take actions in connected services from visuals and alerts.
 - **Generative capabilities - Allow executive summary** to leverage generative capabilities to build precise summary for each sheet of the dashboard.
 - **Allow Data Q&A** to enable generative Q&A for this dashboard. This setting enables Amazon Q in QuickSight.
 - **Allow sharing data stories** to allow viewers of this dashboard to share data stories, that include snapshots of the dashboards visuals, with other users.
 - **Data snapshot options - Allow sharing scenarios** to allow viewers of this dashboard to share scenarios, that include snapshots of the dashboard's visuals, with other users.
6. Choose **Publish**.

The dashboard is now published.

7. On the **Share dashboard** page that opens, choose the **X** icon to close it. You can share the dashboard later by using the sharing option on the dashboard page.

Sample data

To get a first look at how Amazon Quick Sight works, you can explore Amazon Quick Sight using the following sample data:

- [B2b sales data](#)

- [Business overview data](#) (revenue)
- [ML insights data](#)
- [People overview data](#) (human resources)
- [Sales pipeline data](#)
- [Web and social media analytics data](#) (marketing)

Also, a variety of datasets are available free online that you can use with Amazon Quick Sight, for example the [AWS public datasets](#). These datasets come in a variety of formats.

Administering Amazon Quick

The administration experience for Amazon Quick depends on how your account was created. Choose the guide that matches your account type:

- **Free, Plus, and Max accounts** – For accounts created with email or social on aws.com/quick. See [the section called “Free, Plus, and Max accounts”](#).
- **Enterprise accounts** – For accounts provisioned through the AWS Management Console or AWS CLI. These accounts provide governance controls and the full range of enterprise capabilities, including connections to additional enterprise applications, data sources, and knowledge bases. Access is managed through IAM permissions, and customers can connect to Amazon Quick using their enterprise SSO and identity provider. See [the section called “Enterprise accounts”](#).

Administering Amazon Quick (Free/Plus/Max)

This section covers administration for Amazon Quick Free, Plus, and Max accounts that you create at aws.com/quick. Free, Plus, and Max accounts use email or social login for authentication and don't require an AWS account. You can manage users, assets, billing, and account settings directly from the Amazon Quick web interface.

Topics

- [Amazon Quick plans and pricing](#)
- [Upgrading your Amazon Quick plan](#)
- [Managing a Amazon Quick Free or Plus account](#)
- [Managing users in a Amazon Quick Free or Plus account](#)
- [Managing assets in a Amazon Quick Free or Plus account](#)
- [Plan and billing for Amazon Quick Free, Plus, and Max accounts](#)
- [Understanding usage on the Amazon Quick Plus and Max plans](#)
- [Account settings for Amazon Quick Free and Plus accounts](#)
- [User profile and preferences in Amazon Quick](#)

Amazon Quick plans and pricing

Amazon Quick accounts created at aws.com/quick are available in three plans: **Free**, **Plus**, and **Max**. New accounts receive 30 days of Plus at no cost. For current pricing details, plan limits, and feature comparisons, see [Amazon Quick pricing](#).

Note

This page covers pricing for Amazon Quick Free, Plus, and Max accounts created at aws.com/quick. For pricing information about Amazon Quick accessed through the AWS Management Console, see [Amazon Quick pricing](#).

Topics

- [Compare plans](#)
- [Viewing your current plan](#)

Compare plans

Amazon Quick has two account types: accounts created at aws.com/quick (available in the Free, Plus, and Max plans), and Enterprise accounts provisioned through the AWS Management Console, AWS Command Line Interface, or the Amazon Quick API. Most features, including chat with AI agents, spaces, Amazon Quick Flows, Amazon Quick Research, Apps in Amazon Quick, extensions, and integrations, are available to both account types.

Some features are available only to Enterprise accounts. These include Amazon Quick Sight dashboards and analytics, Amazon Quick Automate, and API access.

For a detailed comparison of features, limits, and pricing across all plans, see [Amazon Quick pricing](#).

Viewing your current plan

To view your current plan:

1. In the Amazon Quick navigation bar, choose your username.
2. Choose **Manage account**.
3. In the left navigation, choose **Plan and billing**.

The **Plan & Billing** page displays your current plan name, estimated charges, days remaining (for trial plans), and payment method (for paid plans).

Alternatively, you can choose **Upgrade plan** from the profile menu, which takes you directly to the plan management page.

Upgrading your Amazon Quick plan

You can upgrade your Amazon Quick plan to the Plus or Max plan at any time. This page walks you through the full upgrade process, including entering contact information, verifying your phone number, and adding a payment method.

Note

This page applies to Amazon Quick accounts created with email or social at aws.com/quick. For information about subscription management for Amazon Quick accessed through the AWS Management Console, see [Administering Amazon Quick](#).

Topics

- [Accessing plan management](#)
- [Step-by-step upgrade procedure](#)
- [How upgrade billing works](#)
- [Downgrading your plan](#)

Accessing plan management

You can access the plan management page in two ways.

Option 1: Through the profile menu

1. In the Amazon Quick navigation bar, choose your username.
2. Choose **Upgrade plan**. The plan management page opens in a new browser tab.

Option 2: Through Manage Account

1. In the Amazon Quick navigation bar, choose your username.
2. Choose **Manage account**.

3. In the left navigation of the Manage Account page, choose **Plan and billing**.
4. On the Plan & Billing page, choose **Manage plan**.

Step-by-step upgrade procedure

To upgrade to the Plus or Max plan

1. **Review available plans.**

On the **Manage plan** page, you see the available plans.

Your current plan is indicated with a "Current plan" badge and the number of days remaining (for trial plans).

To see a detailed feature comparison, expand the **Compare plans** section at the bottom of the page.

2. **Select your new plan.**

Under the plan you want to upgrade to, choose **Choose plan**. A **Continue** button appears at the bottom of the page.

Choose **Continue** to proceed.

3. **Enter your contact information.**

You are redirected to a contact information form.

Choose **Continue** after completing all required fields.

4. **Verify your phone number.**

Amazon Quick needs to verify your phone number for security purposes.

Note

If you don't receive the code, choose **Resend code** after the countdown timer completes. Make sure your phone can receive SMS messages from short codes.

- a. On the **Verify your phone number** page, confirm that the displayed phone number is correct. If not, choose **Edit** to change it.

- b. Choose **Send text message**. Amazon Quick sends a 4-digit verification code to your phone via SMS.
 - c. Enter the 4-digit verification code in the **Verification code** field.
 - d. Choose **Continue**.
5. **Enter payment information.**

 Note

You might be redirected to your bank's website to authorize a verification charge. This is a temporary hold and is not an actual charge.

6. **Complete the upgrade.**

After entering your payment information, choose **Upgrade plan**.

Amazon Quick processes your payment method and activates your new plan. You are redirected to the **Plan & Billing** page.

Your upgrade takes effect immediately. All team collaboration features are now available.

How upgrade billing works

When you upgrade your plan, the change takes effect immediately.

 Note

The prices in the following examples are for illustration only. For current rates, see [Amazon Quick plans and pricing](#).

Monthly plan upgrades

- You receive a full credit for your current plan's monthly charge.
- You are then billed the full rate of your new plan for the entire month.
- You receive the full usage allocation of your new plan (agent hours and storage) right away for the entire billing cycle – not prorated for the remaining days.

Example: In this example, assume the Plus plan costs USD \$25 per user per month, and you have 3 users (USD \$75 per month). On August 16, you upgrade. You receive a USD \$75 credit for the Plus charge, and are billed the full rate of your new plan for the entire month. Your usage allocation resets to the new plan's limits immediately.

First-month exception

If you upgrade during your first billing month, both the credit and the new charge are prorated. The proration is based on the days remaining in the month. Starting the following month, the standard full-credit/full-charge model applies.

Example: In this example, assume you subscribe to Plus on July 15, with a prorated charge of approximately USD \$13.71 for 17 remaining days. On July 22, you upgrade. Your credit is USD \$13.71 (the prorated Plus charge), and your new charge is prorated for the same 17-day period. Starting August 1, you are billed the full monthly rate.

Annual plan upgrades

If you are on an annual plan and upgrade mid-term:

- You receive a credit for the unused portion of your annual subscription (the time remaining on your current plan).
- You are charged the new plan's annual rate for the remaining term.
- The time already used on your previous plan is not adjusted.
- You receive the full usage allocation of your new plan immediately.

Example: In this example, assume you are on the annual Plus plan at USD \$240 per year. On day 120 of your term (245 days remaining), you upgrade to a higher annual plan. You receive a credit of approximately USD \$161.10 ($245/365$ of USD \$240) for unused Plus time. You are charged the new plan's annual rate prorated for the remaining 245 days. Your usage allocation increases immediately.

Upgrading from the free trial

- When you upgrade from the Free plan to a paid plan during your free trial, you immediately gain access to the full features and usage allocation of the paid plan.
- Your free trial ends at the point of upgrade and cannot be resumed or restarted.
- Your first charge is prorated for the remaining days in the month. Your next full billing cycle starts on the 1st of the following month.

Example: In this example, assume the Plus plan costs USD \$25 per user per month. You start a free trial on July 1. On July 5, you upgrade to Plus. You are charged approximately USD \$21.77 (USD \$25 multiplied by 27 of 31 days remaining in July). You immediately receive the full Plus allocation. Your free trial cannot be reactivated. Starting August 1, you are billed the full monthly rate.

Downgrading your plan

How to downgrade

To downgrade from Plus or Max to Free:

1. Navigate to **Plan and billing** → **Manage plan**.
2. Under **Free**, choose **Choose plan**.
3. Follow the confirmation prompts.

How downgrade billing works

- Downgrades take effect at the **start of your next billing cycle**, not immediately.
- You continue to have access to your current plan's features and usage allocation through the end of your current billing cycle.
- No refund is issued for the current billing period.
- Starting the next billing cycle, you are billed at the lower plan's rate.

Example: You are on the Plus plan with a billing cycle of August 1–31. On August 10, you request a downgrade to Free. You keep Plus-level access through August 31. Starting September 1, your account reverts to the Free plan.

Annual plan downgrades

- If you are on an annual plan, downgrades take effect at the **end of your annual term**.
- Annual subscriptions cannot be canceled mid-term. There is no refund for the remaining period.

Important

Downgrading to the Free plan reduces your account to a single user and removes access to team collaboration features. All data and configurations are preserved, but only the account owner retains access.

Managing a Amazon Quick Free or Plus account

Amazon Quick Free and Plus accounts – created through the direct signup process at aws.com/quick – include a built-in administration interface for managing your account, users, assets, billing, and support. This interface is accessible directly within the Amazon Quick web application, without requiring access to the AWS Management Console.

Note

This section applies to Amazon Quick Free and Plus accounts only. If your organization accesses Amazon Quick through the AWS Management Console, see [Administering Amazon Quick](#) for information about managing your account using IAM, IAM Identity Center, and other AWS services.

Topics

- [Account administration overview](#)
- [Accessing the Manage Account page](#)
- [Manage Account page layout](#)
- [Free and Plus vs. AWS Console administration](#)

Account administration overview

When you create a Amazon Quick Free or Plus account at aws.com/quick, the person who completes the signup process is automatically assigned the **administrator** role. As the account administrator, you have full access to all account management capabilities, including user management, asset oversight, plan and billing configuration, and support resources.

Free and Plus account administration is organized into the following sections:

Section	Description
Users	Invite team members, manage active users, view usage statistics, and manage pending invitations. For more information, see Managing users in a Amazon Quick Free or Plus account .
Assets	View and manage all account resources including chat agents, spaces, automations, and built-in tools. For more information, see Managing assets in a Amazon Quick Free or Plus account .
Plan and billing	View your current plan, estimated charges, payment methods, and billing history. Upgrade or change your plan. For more information, see Plan and billing for Amazon Quick Free and Plus accounts .
Support	Access AWS Support resources for help with your Amazon Quick account.

Accessing the Manage Account page

To access account administration

1. Sign in to Amazon Quick at aws.com/quick.
2. From the navigation panel, choose your username (displayed with your initials and name).
3. From the profile menu, choose **Manage account**.

The Manage Account page opens, displaying an overview dashboard with a summary of your account sections.

Manage Account page layout

The Manage Account page uses a left-side navigation panel with the following sections:

- **Manage account** – The overview dashboard. Displays a summary card for each section, including a user count and a **Manage** link that navigates directly to the Users page.
- **Users** – Manage active users and pending invitations. See [Managing users](#).
- **Assets** – View and manage account resources. See [Managing assets](#).

- **Plan and billing** – Opens the plan and billing settings page in a new browser tab. See [Plan and billing](#).
- **Support** – Opens the AWS Support portal in a new browser tab for submitting support requests.

Free and Plus vs. AWS Console administration

The following table summarizes the key differences between Free and Plus account administration and AWS Console–based administration.

Feature	Free or Plus account (aws.com/quick)	AWS Console account
Identity provider	Email or social login	IAM Identity Center, IAM, or federated identity
Admin interface	Built-in Quick web UI (Manage Account page)	AWS Management Console
User management	Email-based invitations	IAM Identity Center groups, IAM users
Billing	Billing portal (account.global.app.aws)	AWS Billing and Cost Management console
Plans	Free, Plus	Standard AWS subscription pricing
Account creation	Automatic (AWS account created behind the scenes)	Manual AWS account setup required

Note

Regardless of how your Amazon Quick account was created, all product features – including chat, chat agents, spaces, research, automations, Amazon Quick Sight, integrations, and extensions – function identically. The differences are limited to account setup, identity management, and billing.

Managing users in a Amazon Quick Free or Plus account

In a Amazon Quick Free or Plus account, the account administrator can invite team members by email, monitor user activity, and manage pending invitations. User management is handled directly within the Amazon Quick web application through the Manage Account page.

Note

This section applies to Amazon Quick Free and Plus accounts created at aws.com/quick. For information about managing users in AWS Console–based accounts, see [Managing user access inside Amazon Quick](#).

Topics

- [User limits by plan](#)
- [Navigating to user management](#)
- [Managing active users](#)
- [Inviting new users](#)
- [Managing invited users](#)

User limits by plan

The number of users you can have in your account depends on your plan. For specific user limits for each plan, see [Amazon Quick pricing](#).

For more information about plans, see [Amazon Quick plans and pricing](#).

Navigating to user management

To access the Users page

1. Sign in to Amazon Quick at aws.com/quick.
2. From the navigation panel, choose your username.
3. Choose **Manage account**.
4. In the left navigation, choose **Users**.

The Users page displays two tabs: **Active** and **Invited**.

Managing active users

The **Active** tab displays all users who have accepted their invitation and are currently active in your Amazon Quick account.

Topics

- [Active users table](#)
- [Viewing user usage](#)

Active users table

The Active users table includes the following columns:

- **Name** – The user's display name, as set in their AWS Builder ID profile.
- **Email** – The email address associated with the user's AWS Builder ID.

You can sort the table by either column by choosing the column header. Use the **Search users** field to filter the list by name or email.

Viewing user usage

You can view usage statistics for any active user.

To view a user's usage

1. On the **Active** tab, locate the user in the table.
2. In the row for that user, choose the **More options** menu.
3. Choose **See usage**.

A dialog opens displaying the user's usage statistics.

4. Review the usage information. When finished, choose **Close** to dismiss the dialog.

Inviting new users

You can invite team members to join your Amazon Quick account by sending them an email invitation. Invited users receive an email with instructions to create an AWS Builder ID (if they don't already have one) and join your Quick account.

Note

You can only invite users up to the maximum allowed by your plan. If you need to invite more users than your plan allows, consider [upgrading your plan](#).

To invite new users

1. On the Users page, choose **Invite**.

The Invite dialog opens.

2. In the **Enter email addresses** field, type the email address of the person you want to invite.
3. Press **Enter**. The email address appears as a chip (bubble) in the field, confirming it has been accepted.
4. (Optional) Repeat the previous steps to add additional email addresses. You can also enter multiple email addresses separated by commas.
5. Choose **Send invitations**.

Amazon Quick sends an invitation email to each address. The invitations appear in the **Invited** tab.

Tip

You can invite multiple users at once by entering several email addresses before choosing **Send invitations**. Each email address must be confirmed as a chip in the field before sending.

Managing invited users

The **Invited** tab displays all pending invitations that have not yet been accepted.

Topics

- [Invited users table](#)
- [Resending or canceling invitations](#)

Invited users table

The Invited users table includes the following columns:

- **Email** – The email address the invitation was sent to.
- **Invitation Status** – The current status of the invitation (for example, *Pending*).

You can sort the table by either column. Use the **Search invitations** field to filter the list.

Resending or canceling invitations

You can resend an invitation if the original email was not received, or cancel an invitation to revoke access before it is accepted.

To resend an invitation

1. On the **Invited** tab, locate the invitation in the table.
2. In the row for that invitation, choose the **More options** menu.
3. Choose **Resend**.

Amazon Quick sends a new invitation email to the address.

To cancel an invitation

1. On the **Invited** tab, locate the invitation in the table.
2. In the row for that invitation, choose the **More options** menu.
3. Choose **Cancel**.

The invitation is revoked. The recipient will no longer be able to use the invitation link to join your account.

Important

Canceling an invitation does not affect users who have already accepted and are active in your account. To remove an active user, use the options available on the **Active** tab.

Managing assets in a Amazon Quick Free or Plus account

In a Amazon Quick Free or Plus account, assets represent the resources created and used within your account – including chat agents, spaces, automations, and other AI-powered tools. The Assets page in the Manage Account interface lets administrators view, search, filter, and manage all assets across the account.

Note

This section applies to Amazon Quick Free and Plus accounts created at aws.com/quick. For information about managing resources in AWS Console–based accounts, see [Administering Amazon Quick](#).

Topics

- [Navigating to asset management](#)
- [Viewing assets by name](#)
- [Viewing assets by user](#)
- [Using bulk actions](#)
- [Built-in assets](#)
- [Managing individual assets](#)

Navigating to asset management

To access the Assets page

1. Sign in to Amazon Quick at aws.com/quick.
2. From the navigation panel, choose your username.
3. Choose **Manage account**.
4. In the left navigation, choose **Assets**.

The Assets page opens, displaying all resources in your account.

Viewing assets by name

The default view on the Assets page is **By asset name**, which lists all assets in a single table.

Topics

- [Assets table columns](#)
- [Sorting and searching](#)

Assets table columns

- **Name** – The display name of the asset (for example, *My Assistant* or *RFP Response Generator*).
- **Type** – The asset type (for example, *Chat agent*, *Space*, or *Automation*).

Each row includes:

- A **checkbox** for selecting the asset (used with bulk actions).
- A **More options** menu for asset-specific actions.

Sorting and searching

- Choose a column header to sort the table by that column.
- Use the **Search by asset name** field to filter the list by asset name.
- Use the **All types** dropdown to filter by a specific asset type.

Viewing assets by user

Choose the **By user** tab to view assets grouped by the user who created them.

In this view:

- Use the **Search by user name** combobox to find a specific user and view their assets.
- Use the **All types** dropdown to filter results by asset type.

This view is useful for understanding how assets are distributed across team members and for auditing resource ownership.

Using bulk actions

You can perform actions on multiple assets at once using the bulk action feature.

To use bulk actions

1. On the Assets page, select the checkboxes next to the assets you want to act on. To select all assets, choose the **Select all** checkbox in the table header.

After selecting one or more assets, the **Bulk action** button becomes active.

2. Choose **Bulk action** and select the desired action from the dropdown menu.

Note

The **Bulk action** button is disabled until at least one asset is selected.

Built-in assets

Every new Amazon Quick account includes a set of built-in assets that provide ready-to-use AI-powered tools for common business tasks. These assets are available immediately after account creation and do not count against any resource limits.

Asset name	Description
My Assistant	The default chat agent for your account. Provides general-purpose AI assistance for questions, analysis, and document creation.
Whiteboard Notes Generator	Generates structured notes from whiteboard images or meeting discussions.
Flows Idea Generator	Suggests automation workflow ideas based on your described business processes.
5-Why Root Cause Analysis	Guides you through a structured root cause analysis using the 5-Why methodology.
Marketing Video Analyzer	Analyzes marketing video content and provides insights on messaging, audience, and effectiveness.

Asset name	Description
Social Media Content Creator	Creates social media posts, captions, and content plans tailored to your brand and goals.
Job Description Generator	Drafts professional job descriptions based on role requirements and company context.
Product Launch Email Generator	Creates email campaigns for product launches, including subject lines and body content.
Customer Interest Optimizer	Analyzes customer data to identify interest patterns and suggest optimization strategies.
Flows Prompt Helper	Assists with writing effective prompts for Amazon Quick automations and flows.
RFP Response Generator	Drafts responses to Requests for Proposal (RFPs) based on your organization's capabilities and the RFP requirements.

 Tip

Built-in assets serve as starting points. You can customize these assets or create new ones tailored to your organization's specific needs. For more information about creating custom chat agents, see [Create and deploy AI assistants using Amazon Quick chat agents](#).

Managing individual assets

To manage a specific asset, choose the **More options** menu in the asset's row. The available actions depend on the asset type and may include options such as viewing details, editing, sharing, or deleting the asset.

Plan and billing for Amazon Quick Free, Plus, and Max accounts

This topic describes how to view and manage your plan and billing information for Amazon Quick Free, Plus, and Max accounts created at aws.com/quick.

 Note

This section applies to Amazon Quick Free, Plus, and Max accounts. If your organization manages Amazon Quick through the AWS Management Console, billing is handled through your AWS account. For more information, see [AWS Billing and Cost Management](#).

Topics

- [Accessing Plan & Billing](#)
- [Viewing your current plan](#)
- [Billing history](#)
- [Billing preferences](#)
- [How collaborator billing works](#)
- [Annual subscription terms](#)
- [Billing cycle and proration](#)
- [Changing your plan](#)

Accessing Plan & Billing

The Plan & Billing page is located in the Amazon Quick account settings portal.

To access Plan & Billing

1. In Amazon Quick, from the navigation bar, choose your username.
2. Choose **Manage account**.
3. In the left navigation of the Manage Account page, choose **Plan and billing**.

The Plan & Billing settings page opens in a new browser tab at `account.global.app.aws`.

The settings page has its own navigation panel with the following sections:

- **Plan & Billing** – View your current plan, estimated charges, and billing history.
- **My account** – View your account profile and account details. For more information, see [Account settings for Amazon Quick Free and Plus accounts](#).
- **Support** – Access AWS Support (opens externally).

Viewing your current plan

The Plan & Billing page displays the following information about your current subscription:

- **Current plan** – The name and type of your active plan. For example, **FREE**, **PLUS**, or **MAX**.
- **Manage plan** – Choose this button to view available plans and change your subscription. For more information, see [Upgrading your Amazon Quick plan](#).
- **Current estimated charges** – The estimated cost for your current billing period, displayed in USD.
- **Days remaining** – For trial plans, the number of days remaining in the trial period along with the expiration date.
- **Payment method** – The card type and last four digits of the payment method on file. A **Change** link is available to update your payment method.
- **Estimated charges breakdown** – An expandable section that provides a detailed breakdown of the charges for the current billing period.

Billing history

The **Billing history** section displays a table of past billing records with the following columns:

Column	Description
ID	A unique identifier for the billing record.
Type	The type of charge (for example, subscription).
Status	The payment status (for example, paid, pending).
Due date	The date the payment is due.
Payment method	The payment method used for the charge.
Amount	The total amount charged.
Charge type	The category of the charge.
Action	Available actions for the billing record, such as downloading an invoice.

 Note

New accounts will not have any billing records until the first billing cycle completes. The table displays "No billing records to display" until charges are processed.

Billing preferences

Choose the **Preferences** button on the Plan & Billing page to configure your billing notification preferences.

How collaborator billing works

When you add a collaborator to your Amazon Quick Plus or Max account, billing for their seat begins when the collaborator **accepts the invitation** – not when the invitation is sent.

When billing starts

- You incur no charges for a pending invitation. You are charged only for collaborators who accept.
- When a collaborator accepts, billing begins for their seat at a prorated rate for the remainder of your current billing cycle.
- The prorated charge appears on your next invoice. It is not charged to your payment method on the day the invitation is accepted.

Removing a collaborator

- Removing a collaborator does not result in a refund for the current billing period.
- The seat charge applies through the end of your current billing cycle, regardless of when the collaborator is removed.
- This applies even if the collaborator is removed shortly after accepting the invitation.

Canceling a pending invitation

- If a collaborator never accepts the invitation, your account incurs no charge for that seat.
- You can cancel a pending invitation at any time at no cost.

Annual subscription terms

When you commit to an annual plan, your subscription is billed for the full annual term.

- Annual subscriptions cannot be canceled mid-term.
- Cancellations and downgrades take effect at the end of your annual term. You receive no refund for the remaining period.
- Upgrades to a higher plan during your annual term apply a credit for the unused portion of your current annual subscription, and you are charged the new plan's annual rate for the remainder of your term.

Billing cycle and proration

Amazon Quick bills on a calendar-month basis.

- Your billing month starts at 12:00 AM (UTC) on the first day of the month and ends at 11:59 PM (UTC) on the last day.
- If you subscribe partway through a month, your first charge is prorated for the remaining days in that month. Your next full billing cycle starts on the 1st of the following month.
- If you upgrade mid-month, you are billed immediately for the full new plan rate minus a full credit for your previous plan's monthly charge. Overage charges, if applicable, appear as a separate line item at the end of the billing cycle.

Changing your plan

To upgrade, downgrade, or change your Amazon Quick plan, choose **Manage plan** on the Plan & Billing page. For detailed instructions on the upgrade process, including entering contact information and payment details, see [Upgrading your Amazon Quick plan](#).

For information about the available plans and their features, see [Amazon Quick plans and pricing](#).

Understanding usage on the Amazon Quick Plus and Max plans

Amazon Quick uses limits to keep your usage predictable and to spread heavy activity evenly across your billing month. Understanding how these limits work helps you avoid interruptions, decide when to enable more usage, and get the most from your plan.

 Note

This page applies to Amazon Quick Free, Plus, and Max accounts created with email or social login at aws.com/quick. It does not apply to Amazon Quick Professional or Enterprise plans managed through the AWS Management Console, which handle usage differently.

Topics

- [How usage limits work](#)
- [Usage does not roll over](#)
- [How to get more usage](#)
- [Best practices to conserve usage](#)

How usage limits work

Your usage is governed by two limits that work together: a monthly limit and a session limit.

- **Monthly limit** – The total amount of work you can do with Quick in a billing month. This limit resets at the start of each billing cycle.
- **Session limit** – A shorter, rolling window that limits activity in any single stretch. The session limit is a percentage of your monthly limit. It refreshes continuously as time passes, so you no longer need to wait until the next billing month. When you reach a session limit, Quick pauses your session briefly. It does not block you for the rest of the month.

 Note

Reaching a session limit is temporary. It clears on its own as the rolling window moves forward, so you can keep working after a short pause.

Quick measures your usage by the length and complexity of your work, the features you use, the model selected, and the thinking level. Heavier tasks use more; lighter tasks use less. Because usage depends on these factors, the same activity can consume different amounts on different days.

What to do when you hit a limit

Limit	Free plan	Plus and Max plans
Session limit	Wait for the session window to refresh, or upgrade to Plus or Max for more usage.	Wait for the session window to refresh, or enable additional usage to keep working.
Monthly limit	Upgrade to Plus or Max, or wait for the reset at the start of the next billing cycle.	Enable additional usage, or wait for the reset at the start of the next billing cycle.

Usage does not roll over

Your usage allocation resets at the start of each billing cycle. Unused usage from the previous cycle does not carry over, and no credit is issued for unused capacity. Make the most of your allocation each month.

How to get more usage

If you need more usage, you can enable additional usage and set a maximum monthly limit—a spend cap that you control.

- **You control the cost** – You only pay for what you actually use, up to the cap you set. The cap is a ceiling, not a charge.
- **What changes** – Once enabled, your monthly capacity expands up to your cap, and your session windows adjust accordingly, so you're less likely to be paused mid-task.

To enable additional usage

1. In Amazon Quick, from the navigation bar, choose your username, and then choose **Manage account**.
2. Enable additional usage.
3. Set your maximum monthly limit (spend cap).

How additional usage billing works

- Additional usage is **off by default**. You must manually enable it in your account settings.
- When you enable additional usage, you are required to set a **maximum monthly limit** (spending cap). This prevents unexpected charges.
- When additional usage is enabled and your cap has not been reached, your work continues uninterrupted beyond your plan's included allocation.
- Additional usage is billed at the end of your billing cycle as a separate line item on your invoice.
- **When your spending cap is reached:** Usage pauses until your next billing cycle resets your allocation. You can also raise your cap or upgrade your plan to continue.

What happens when you upgrade with additional usage charges

If you have accumulated additional usage charges and then upgrade to a higher plan mid-cycle, your usage is **recalculated against the new plan's higher allocation**.

- The higher plan includes more capacity. If your total usage for the month now falls within the new plan's included allocation, your additional usage charges are reduced or eliminated.
- You are only billed for usage that exceeds the new plan's allocation.

Example: You are on the Plus plan and have used more than your plan's included allocation this month, so you have additional usage charges pending. You upgrade to a higher plan with a larger included allocation. Your month-to-date usage now falls within the new plan's allocation, so your additional usage charges are eliminated.

Tip

If your additional usage charges are approaching the price difference between your current plan and a higher plan, upgrading might lower your total bill and increase your included allocation.

Additional usage and downgrades

- If you downgrade your plan, your current allocation remains in effect through the end of your billing cycle.

- Downgrading does not recalculate your additional usage charges for the current cycle.
- Your new, lower allocation takes effect at the start of your next billing cycle.

Best practices to conserve usage

Use the following practices to make your usage go further before you need to enable additional usage:

- **Start fresh for new tasks** – Start a new chat for a new topic. Long, growing conversations carry their whole history with each turn and consume more usage. A fresh chat is lighter.
- **Match the thinking level to the task** – Use a lower thinking level for routine tasks, and reserve the highest thinking level for genuinely complex work. Higher thinking levels consume more usage.
- **Be concise and specific** – Clear, scoped prompts avoid long back-and-forth and re-work, which saves usage.

Account settings for Amazon Quick Free and Plus accounts

This topic describes how to view and manage account-level settings for Amazon Quick Free and Plus accounts created at aws.com/quick. Account settings include your profile information, account details, and support access.

Note

This section applies to Amazon Quick Free and Plus accounts. If your organization manages Amazon Quick through the AWS Management Console, account settings are managed through your AWS account and IAM Identity Center. For more information, see [Administering Amazon Quick](#).

Topics

- [Accessing account settings](#)
- [My account](#)
- [Support](#)

Accessing account settings

The account settings are located in the Amazon Quick account settings portal, alongside Plan & Billing.

To access account settings

1. In Amazon Quick, from the navigation bar, choose your username.
2. Choose **Manage account**.
3. In the left navigation of the Manage Account page, choose **Plan and billing**.

The settings portal opens in a new browser tab at `account.global.app.aws`.

4. In the left navigation of the settings portal, choose **My account**.

My account

The My account page displays your profile information and account details.

Topics

- [Profile & security](#)
- [Account details](#)
- [Account closure and reopening](#)

Profile & security

The Profile & security section displays the following information:

- **Name** – Your display name. This is the name associated with your AWS Builder ID.
- **Email** – Your email address. This is the email associated with your AWS Builder ID.
- **Manage profile & security** – Choose this link to open the AWS Builder ID profile management page in a new browser tab. From there, you can update your name, nickname, email, password, multi-factor authentication settings, and privacy preferences.

Note

Your name and email address in Amazon Quick are managed through your AWS Builder ID. To change these details, choose **Manage profile & security** to open the Builder ID profile

page, or choose **Edit profile** from the Amazon Quick profile menu. For more information, see [User profile and preferences in Amazon Quick](#).

Account details

The Account details section displays the following information:

- **Account ID** – A unique numeric identifier for your Amazon Quick account. You may need this ID when contacting AWS Support.
- **Copy Account ID** – Choose this button to copy your Account ID to the clipboard.
- **Close account** – Choose this button to close your Amazon Quick account. For more information, see [Account closure and reopening](#).

Account closure and reopening

You can close your Amazon Quick account at any time. After closure, you have 90 days to reopen it before it is permanently closed and AWS deletes your content.

To close your account

- Go to **Manage account** → **Plan & Billing** → **My account** → **Close account**.

To reopen your account, contact AWS Support within the 90-day post-closure period.

Before closing, note the following:

- All content and resources are permanently deleted after the 90-day post-closure period, including chat history, agents, spaces, automations, and uploaded files.
- All users lose access immediately.
- You remain responsible for outstanding charges and will receive a final bill the following month.

Support

To access AWS Support for Amazon Quick:

- In the left navigation of the settings portal, choose **Support**.

The AWS Support portal opens in a new browser tab at `support.aws.amazon.com`.

From the AWS Support portal, you can:

- Create a new support case for Amazon Quick.
- View the status of existing support cases.
- Access self-service troubleshooting resources.

 Tip

You can also access support resources from within Amazon Quick by choosing your username, then choosing **Help** in the profile menu. The Help submenu includes links to Getting Started guides, Quick Documentation, the Quick Community forum, and How to Videos.

User profile and preferences in Amazon Quick

This topic describes how to access and configure your personal profile settings, preferences, and related options in Amazon Quick. These settings apply to your individual user experience and are separate from account-level administration settings.

Topics

- [Accessing the profile menu](#)
- [Edit profile](#)
- [Preferences](#)
- [Help](#)
- [Quick anywhere](#)
- [Private Mode](#)

Accessing the profile menu

To access the profile menu, from the Amazon Quick navigation bar, choose your name or avatar. The profile menu displays your name, email address, and current plan information.

The profile menu contains the following options:

Menu option	Description
Upgrade plan	Opens the plan management page in a new browser tab where you can change your subscription tier.
Manage account	Opens the Manage Account page within Amazon Quick. Available to account owners and administrators. For more information, see Managing a Amazon Quick Free or Plus account .
Switch accounts	Allows you to switch between multiple Amazon Quick accounts if you are a member of more than one.
Edit profile	Opens the AWS Builder ID profile management page in a new browser tab.
Preferences	Opens the Preferences dialog where you can configure personalization, chat, and language settings.
Help	Expands a submenu with links to help resources.
Quick anywhere	Expands a submenu with options for accessing Quick on other platforms.
Sign out	Signs you out of your Amazon Quick account.

Edit profile

When you choose **Edit profile** from the profile menu, the AWS Builder ID profile management page opens in a new browser tab. Your AWS Builder ID is the identity you use to sign in to your Amazon Quick Free or Plus account. From this page, you can manage your personal information, security settings, and privacy preferences.

Topics

- [My details](#)
- [Security](#)
- [Privacy & data](#)

My details

The My details tab displays and allows you to edit your personal information.

- **Name** – Your full name. Choose **Edit** to update.
- **Nickname** – An optional display name. Choose **Edit** to update.
- **Email address** – The email address associated with your AWS Builder ID. Choose **Edit** to change your email address. A verification step is required for email changes.

Security

The Security tab provides options for managing your account security.

- **Change password** – Choose this button to update your AWS Builder ID password.
- **Multi-factor authentication (MFA)** – Manage MFA devices for additional sign-in security. The MFA devices table displays the following columns: Name, Type, Make and Model, and Registered Date.
 - **Register device** – Add a new MFA device (authenticator app or security key).
 - **Rename** – Change the display name of a registered MFA device.
 - **Delete** – Remove a registered MFA device.
- **Active sessions** – View all currently active sign-in sessions.
 - The sessions table displays the following columns: Device, IP address, and Sign-in date.
 - **Delete all active sessions** – Choose this button to immediately terminate all active sessions. You will be signed out of all devices.

Privacy & data

The Privacy & data tab provides options for managing your data and privacy.

- **AWS Privacy Notice** – A link to the AWS Privacy Notice for information about how AWS handles your data.
- **Request your Builder ID data** – Choose this button to request a copy of the data associated with your AWS Builder ID.
- **Request your AWS data** – A link to request a copy of all your AWS data through the AWS Data Subject Access Request process.
- **Delete AWS Builder ID** – Choose this button to permanently delete your AWS Builder ID.

Warning

Deleting your AWS Builder ID is permanent and cannot be undone. You will lose access to all AWS services and applications that use your Builder ID for authentication, including Amazon Quick.

Preferences

When you choose **Preferences** from the profile menu, a dialog opens with three tabs for configuring your personal Amazon Quick experience. Choose **Save** to apply changes or **Cancel** to discard them.

Topics

- [Personalization](#)
- [Chat](#)
- [Language and time](#)

Personalization

The Personalization tab allows you to customize how Amazon Quick tailors its responses to your needs.

- **Your current work focus** – A text field where you can describe what you are currently working on (for example, "Q3 financial planning" or "Product launch preparation"). Amazon Quick uses this context to provide more relevant suggestions and responses.
- **Manage memories** – Choose this button to view and manage the information that Amazon Quick has learned about you through your conversations. You can review, edit, or delete individual memories.

Chat

The Chat tab allows you to configure default behaviors for the Amazon Quick chat experience.

- **Assistant data access** – A toggle that controls whether My Assistant can access all your connected data sources by default. When enabled, My Assistant can query your connected

data, documents, and integrations without you needing to specify the data source for each conversation. You can override this setting on a per-conversation basis.

- **Choose default data** – A dropdown selector that sets the default data scope for new conversations. You can override the default data scope within any individual conversation by using the data source selector in the chat footer. Options include:
 - **All data** – My Assistant can access all connected data sources and applications.
 - Other options may appear based on your connected integrations and data sources.

Language and time

The Language and time tab allows you to configure regional and display preferences.

- **Timezone** – Choose your preferred timezone from the dropdown selector. This setting affects how dates and times are displayed throughout Amazon Quick.
- **Language** – Choose your preferred display language from the dropdown selector. This setting controls the language of the Amazon Quick user interface.
- **Use browser time zone** – Select this checkbox to automatically use your web browser's detected timezone instead of the manually selected timezone above. When enabled, the timezone selector is overridden by your browser setting.

Help

When you choose **Help** from the profile menu, a submenu expands with the following resources:

Resource	Description
Getting Started	Opens an interactive guide that walks you through the key features of Amazon Quick.
Quick Documentation	Opens the official Amazon Quick documentation at docs.aws.amazon.com.
Ask the Quick Community	Opens the Amazon Quick community forum where you can ask questions and share knowledge with other users.
How to Videos	Opens a collection of tutorial videos demonstrating Amazon Quick features and workflows.

Quick anywhere

When you choose **Quick anywhere** from the profile menu, a submenu expands with options for accessing Amazon Quick on other platforms:

Platform	Description
Desktop	Download the Amazon Quick desktop application for macOS or Windows. The desktop app provides a native experience with system-level integrations.
Mobile	Access Amazon Quick on your mobile device.
Browser extensions	Install browser extensions that integrate Amazon Quick directly into your web browsing experience.

Private Mode

Amazon Quick includes a **Private Mode** toggle, accessible from the top area of the navigation bar. When Private Mode is enabled:

- Your conversations are not saved to your chat history.
- Amazon Quick does not learn from or retain information from private conversations.
- Private conversations do not appear in the Recents section of the navigation.

To enable or disable Private Mode, choose the **Private Mode** button in the navigation bar.

Administration for Amazon Quick

As an administrator, you decide who can access your Amazon Quick account and what those users can do in it. You also configure the security and encryption settings that protect your data. You do this work in the Amazon Quick administration console. Some sections require AWS Identity and Access Management (IAM) permissions, and you can perform some tasks programmatically with the Amazon Quick API.

If you created your account with email or social on aws.com/quick, see [Administering Amazon Quick \(Free/Plus/Max\)](#) instead.

Use this section to do the following:

- Set up identity and access, including IAM Identity Center, IAM federation, Active Directory, external identity providers, and multi-factor authentication.
- Manage users, groups, and subscriptions.
- Provision SPICE capacity and index capacity for unstructured data.
- Connect to your data through VPC connections and AWS data sources.
- Secure your account with domain allowlisting, namespaces, IP and VPC restrictions, and AWS Key Management Service customer managed keys.
- Control permissions with default access settings, custom permissions, IAM policy assignments, and limit profiles.
- Govern usage with approval workflows and data loss prevention.
- Customize branding and agent settings, and monitor usage metrics.

Topics

- [Getting started with Amazon Quick](#)
- [Setting up for Quick](#)
- [Configuring VPC connections in Amazon Quick Sight](#)
- [Configuring Amazon Quick Sight access to AWS data sources](#)
- [Amazon Quick account management](#)
- [Managing Quick subscriptions](#)
- [Managing identities](#)
- [Managing security](#)
- [Permissions](#)
- [Governance](#)
- [Customizations](#)
- [Using the Amazon Quick analytics dashboard](#)
- [AWS security in Quick](#)

Getting started with Amazon Quick

This chapter provides you an overview of the key Amazon Quick concepts you need to understand before you dive in and start using Quick.

Topics

- [Amazon Quick user types](#)
- [Plan your Quick account structure and AWS Regions](#)
- [AWS Regions, websites, IP address ranges, and endpoints](#)
- [Supported browsers](#)

Amazon Quick user types

Amazon Quick has three distinct user personas based on permission levels and use cases: administrator, author, and reader. Within these personas, there are six specific user roles that can be assigned, including Pro versions that provide access to advanced Amazon Quick capabilities. The following sections describe the role and task boundaries of each type.

Topics

- [Understanding Amazon Quick subscriptions and roles](#)
- [Understanding reader capabilities](#)
- [Understanding author capabilities](#)
- [Understanding admin capabilities](#)

Understanding Amazon Quick subscriptions and roles

The new capabilities for Amazon Quick are available via the following monthly user subscriptions:

- **Amazon Quick Professional** (or Amazon Quick Sight Reader Pro role)
- **Amazon Quick Enterprise** (or Amazon Quick Sight Author Pro role)

To manage Amazon Quick users, billing and services, you need the Admin Pro role, which is the same as Amazon Quick Enterprise user subscription.

Amazon Quick also provides the following monthly user subscriptions to support business only (BI) only use cases:

- **Amazon Quick Sight Reader** – For users who need to view and interact with dashboards and reports.
- **Amazon Quick Sight Author** – For users who need to create and manage BI content.

To manage Amazon Quick Sight users, billing and services, you need Admin role, which is the same as Amazon Quick Sight Author user subscription.

Note

For pricing information, see [Amazon Quick pricing](#)

Understanding reader capabilities

Readers can use Quick to access company data and find answers through chat interactions with AI agents. They can upload files, run automations, create visualizations, and share spaces via direct links. Reader Pro users additionally have access to advanced Amazon Quick features including AI chat agents, collaborative spaces, flows, and extensions. For detailed information about reader capabilities and subscription types, see [Amazon Quick pricing](#).

Understanding author capabilities

Authors are domain experts who build and manage Quick resources. In addition to reader capabilities, they can create datasets, dashboards, automations, and agents. They have expanded sharing permissions for spaces and can use AI to create visualizations. Author Pro users can additionally create content using natural language, build knowledge bases, configure actions, and access advanced automation capabilities. For detailed information about author capabilities and subscription types, see [Amazon Quick pricing](#).

Understanding admin capabilities

Administrators manage user access, monitor costs, and maintain data sources with full reader and author capabilities. There are two types of administrators: system administrators oversee the broader AWS environment including Amazon Quick signup, system health, and security; Amazon Quick administrators manage users and resources within Amazon Quick.

When your Amazon Quick account is integrated with IAM Identity Center, admin capabilities are divided between IAM permissions and Amazon Quick admin role permissions. Access to some sections of the Amazon Quick administration console is restricted by IAM permissions. The

following table summarizes the admin actions that you can perform in Amazon Quick based on the access type that you choose.

Admin action	IAM permissions	Amazon Quick admin role permissions
Manage assets	✔ Yes	✘ No
Security & permissions	✔ Yes	✘ No
Manage VPC connections	✔ Yes	✘ No
KMS keys	✔ Yes	✘ No
Account settings	✔ Yes	✘ No
Account customization	✘ No	✔ Yes
Manage users	✔ Yes (IAM Identity Center users)	✔ Yes (Amazon Quick and IAM users)
Your subscriptions	✘ No	✔ Yes
Mobile settings	✘ No	✔ Yes
Domains and embedding	✘ No	✔ Yes
SPICE capacity	✘ No	✔ Yes

Admin and Admin Pro users have full reader and author capabilities but focus primarily on system administration to ensure efficient and secure operations for all users. For detailed information about author capabilities and subscription types, see [Amazon Quick pricing](#).

Plan your Quick account structure and AWS Regions

This topic describes the recommended way to structure Amazon Quick for your organization, how a single account operates across AWS Regions, and the role of your home Region.

Topics

- [Recommended account structure](#)

- [How a single account works across AWS Regions](#)
- [Your home Region](#)

Recommended account structure

We recommend that you use a single Amazon Quick account for your organization. A single account gives you one place to connect all of your apps and data, one place to manage users, and one place to manage fine-grained permissions. This makes it easier to apply consistent governance, share assets, and audit access across your organization.

To keep teams, projects, or business units separate within a single account, use namespaces. Namespaces provide logical isolation of users and resources while you maintain centralized account management. For more information, see [Supporting multitenancy with isolated namespaces](#).

To keep Quick subscription creation centralized on your approved account, you can use service control policies in AWS Organizations. For more information, see [Using service control policies to restrict Amazon Quick sign-up options](#).

How a single account works across AWS Regions

Amazon Quick is available in multiple AWS Regions. A single Amazon Quick account spans all available Amazon Quick Regions. After you sign up, you can create and use assets in any available Amazon Quick Region, and those assets live in the Region where you create them. These assets include datasets, knowledge bases, agents, spaces, flows, and dashboards. Each AWS account can have one Amazon Quick subscription, and that subscription can be used in multiple AWS Regions.

Capacity is managed separately in each Region. SPICE capacity and Quick Index capacity are allocated per AWS Region, so you provision capacity in each Region where you work with data.

Your home Region

Your home Region is the AWS Region that you select when you set up your Amazon Quick subscription. It is also called your account's identity Region, the Region that anchors your Amazon Quick account identity and user sign-in. Your home Region must be a supported Amazon Quick Region.

Your home Region also sets your default capacity behavior. Amazon Quick automatically creates your index in your home Region with auto-scaling, and your Index capacity allocation is billed against your home Region. If you provision Index capacity in Regions beyond your home Region, that capacity is billed as overage.

For the list of supported Regions, and for how Amazon Q in Quick processes AI inference across Regions within your geography, see [AWS Regions, websites, IP address ranges, and endpoints](#) and [Cross-Region inference for Australia, Japan, Europe, and the United States](#).

AWS Regions, websites, IP address ranges, and endpoints

AWS cloud-computing resources are housed in highly available facilities in different areas of the world (for example, North America, Europe, and Asia). These facilities are each part of an AWS Region. For more information about AWS Regions and Availability Zones (AZs), see [Global infrastructure](#).

The IP addresses listed in the sections below are the ranges where Amazon Quick traffic originates from when making outbound connections to databases. They are not the IP address ranges that you use to connect to the Amazon Quick website or service API. For more information about authorizing Amazon Quick, [Authorizing connections to AWS data sources](#).

Topics

- [Supported AWS Regions for Amazon Quick](#)
- [Cross-Region calls for web search](#)
- [Cross-Region inference for Australia, Japan, Europe, and the United States](#)
- [AWS Regions for Amazon Q in Quick](#)

Supported AWS Regions for Amazon Quick

Amazon Quick is currently supported in the following AWS Regions. The following list provides websites, IP address ranges, and endpoints for Amazon Quick in each AWS Region.

Region name	Region code	Agentic Features	Website for user access	API endpoints (HTTPS)	IP address range for data source connectivity
US East (Ohio)	us-east-2	No	https://us-east-2.quick sigh	quicksigh t.us-east	52.15.247 .160/27

Region name	Region code	Agentic Features	Website for user access	API endpoints (HTTPS)	IP address range for data source connectivity
			t.aws.amazon.com	-2.amazonaws.com	
US East (N. Virginia)	us-east-1	Yes	https://us-east-1.quicksight.aws.amazon.com	quicksight.us-east-1.amazonaws.com	52.23.63.224/27
US West (Oregon)	us-west-2	Yes	https://us-west-2.quicksight.aws.amazon.com	quicksight.us-west-2.amazonaws.com	54.70.204.128/27
Africa (Cape Town)	af-south-1	No	https://af-south-1.quicksight.aws.amazon.com	quicksight.af-south-1.amazonaws.com	13.246.220.192/27
Asia Pacific (Jakarta)	ap-southeast-3	No	https://ap-southeast-3.quicksight.aws.amazon.com	quicksight.ap-southeast-3.amazonaws.com	43.218.71.192/27

Region name	Region code	Agentic Features	Website for user access	API endpoints (HTTPS)	IP address range for data source connectivity
Asia Pacific (Malaysia)	ap-southeast-5	No	<code>https://ap-southeast-5.quicksight.aws.amazon.com</code>	<code>quicksight.ap-southeast-5.amazonaws.com</code>	56.68.33.0/27
Asia Pacific (Mumbai)	ap-south-1	No	<code>https://ap-south-1.quicksight.aws.amazon.com</code>	<code>quicksight.ap-south-1.amazonaws.com</code>	52.66.193.64/27
Asia Pacific (Seoul)	ap-northeast-2	No	<code>https://ap-northeast-2.quicksight.aws.amazon.com</code>	<code>quicksight.ap-northeast-2.amazonaws.com</code>	13.124.145.32/27
Asia Pacific (Singapore)	ap-southeast-1	No	<code>https://ap-southeast-1.quicksight.aws.amazon.com</code>	<code>quicksight.ap-southeast-1.amazonaws.com</code>	13.229.254.0/27

Region name	Region code	Agentic Features	Website for user access	API endpoints (HTTPS)	IP address range for data source connectivity
Asia Pacific (Sydney)	ap-southeast-2	Yes	https://ap-southeast-2.quicksight.aws.amazon.com	quicksight.ap-southeast-2.amazonaws.com	54.153.249.96/27
Asia Pacific (Tokyo)	ap-northeast-1	Yes	https://ap-northeast-1.quicksight.aws.amazon.com	quicksight.ap-northeast-1.amazonaws.com	13.113.244.32/27
Canada (Central)	ca-central-1	No	https://ca-central-1.quicksight.aws.amazon.com	quicksight.ca-central-1.amazonaws.com	15.223.73.0/27
China (Beijing)	cn-north-1	No	https://cn-north-1.quicksight.amazonaws.cn	quicksight.cn-north-1.amazonaws.com.cn	71.136.65.64/27

Region name	Region code	Agentic Features	Website for user access	API endpoints (HTTPS)	IP address range for data source connectivity
Europe (Frankfurt)	eu-central-1	Yes	https://eu-central-1.quicksight.aws.amazon.com	quicksight.eu-central-1.amazonaws.com	35.158.127.192/27
Europe (Ireland)	eu-west-1	Yes	https://eu-west-1.quicksight.aws.amazon.com	quicksight.eu-west-1.amazonaws.com	52.210.255.224/27
Europe (London)	eu-west-2	Yes	https://eu-west-2.quicksight.aws.amazon.com	quicksight.eu-west-2.amazonaws.com	35.177.218.0/27
Europe (Milan)	eu-south-1	No	https://eu-south-1.quicksight.aws.amazon.com	quicksight.eu-south-1.amazonaws.com	18.102.150.128/27

Region name	Region code	Agentic Features	Website for user access	API endpoints (HTTPS)	IP address range for data source connectivity
Europe (Paris)	eu-west-3	No	<code>https://eu-west-3.quickights.aws.amazon.com</code>	<code>quicksight.eu-west-3.amazonaws.com</code>	<code>13.38.202.0/27</code>
Europe (Spain)	eu-south-2	No	<code>https://eu-south-2.quickights.aws.amazon.com</code>	<code>quicksight.eu-south-2.amazonaws.com</code>	<code>18.101.99.160/27</code>
Europe (Stockholm)	eu-north-1	No	<code>https://eu-north-1.quickights.aws.amazon.com</code>	<code>quicksight.eu-north-1.amazonaws.com</code>	<code>13.53.191.64/27</code>
Europe (Zurich)	eu-central-2	No	<code>https://eu-central-2.quickights.aws.amazon.com</code>	<code>quicksight.eu-central-2.amazonaws.com</code>	<code>16.63.53.32/27</code>

Region name	Region code	Agentic Features	Website for user access	API endpoints (HTTPS)	IP address range for data source connectivity
South America (São Paulo)	sa-east-1	No	https://sa-east-1.quickstart.aws.amazon.com	quickstart.us-east-1.amazonaws.com	18.230.46.192/27
AWS GovCloud (US-East)	gov-east-1	No	https://quickstart.us-gov-east-1.amazonaws.com	quickstart.us-gov-east-1.amazonaws.com	18.252.16.5.64/27
AWS GovCloud (US-West)	gov-west-1	Yes	https://quickstart.us-gov-west-1.amazonaws.com	quickstart.us-gov-west-1.amazonaws.com	160.1.180.32/27
Israel (Tel Aviv)	il-central-1	No	https://il-central-1.quickstart.aws.amazon.com	quickstart.il-central-1.amazonaws.com	51.17.195.32/27

Region name	Region code	Agentic Features	Website for user access	API endpoints (HTTPS)	IP address range for data source connectivity
Middle East (UAE)	me-central-1	No	https://me-central-1.quicksight.aws.amazon.com	quicksight.me-central-1.amazonaws.com	51.112.11.224/27

If your IAM Identity Center instance is replicated to multiple Regions, you can create your Quick Enterprise subscription in an active additional Region. For more information, see [Set up Amazon Quick with IAM Identity Center multi-Region](#).

Cross-Region calls for web search

Amazon Quick makes cross-Region calls for web search functionality in Chat, Agents, and Research features. When using web search, requests originating in US West (Oregon) are processed in the US East (N. Virginia) Region. For customers in EMEA Regions, web search queries are processed in the Europe (Ireland) Region.

Web search regional availability

Region name and code	Web search processing region
US East (N. Virginia) (us-east-1)	US East (N. Virginia) (us-east-1)
US West (Oregon) (us-west-2)	US East (N. Virginia) (us-east-1)
Europe (Ireland) (eu-west-1)	Europe (Ireland) (eu-west-1)
Europe (Frankfurt) (eu-central-1)	Europe (Ireland) (eu-west-1)

Region name and code	Web search processing region
Europe (London) (eu-west-2)	Europe (Ireland) (eu-west-1)
Asia Pacific (Sydney) (ap-southeast-2)	Asia Pacific (Sydney) (ap-southeast-2)
Asia Pacific (Tokyo) (ap-northeast-1)	Asia Pacific (Tokyo) (ap-northeast-1)

Cross-Region inference for Australia, Japan, Europe, and the United States

With cross-Region inference, Amazon Quick automatically selects the optimal Region within your geography to process your inference request. This approach maximizes available compute resources and model availability. With cross-Region inference, you get:

- Complete access to the most advanced Amazon Quick AI capabilities and features
- Access to a variety of models suitable for different tasks
- Improved performance for all your applications

Cross-Region inference requests stay within the AWS Regions in the geography where your data originally resides. For example, a request made in the US stays within the AWS Regions in the US. Although data is stored only in your primary Region, your input prompts and output results might move outside of that Region when you use cross-Region inference. Amazon transmits all data encrypted across its secure network.

Note

There is no additional cost for using cross-Region inference. Amazon CloudWatch and AWS CloudTrail logs do not specify the AWS Region in which data inference occurs.

Supported Regions for cross-Region inference

For a list of Region codes and endpoints supported in Amazon Quick AI capabilities and features, see [Supported AWS Regions for Amazon Quick](#).

The following table summarizes the supported geographies and their inference Regions for Amazon Quick cross-Region inference.

Supported Amazon Quick geography	Inference Regions
United States	US East (N. Virginia) (us-east-1)
	US East (Ohio) (us-east-2)
	US West (Oregon) (us-west-2)
Europe	Europe (Frankfurt) (eu-central-1)
	Europe (Ireland) (eu-west-1)
	Europe (Paris) (eu-west-3)
	Europe (Stockholm) (eu-north-1)
	Europe (Milan) (eu-south-1)
	Europe (Spain) (eu-south-2)
Australia	Asia Pacific (Sydney) (ap-southeast-2)
	Asia Pacific (Melbourne) (ap-southeast-4)
Japan	Asia Pacific (Tokyo) (ap-northeast-1)
	Asia Pacific (Osaka) (ap-northeast-3)

Note

Amazon Quick follows the Amazon Bedrock Europe geographic cross-Region inference profile. For the Europe geography:

- Amazon Quick processes a request that originates in a Europe Region other than Europe (London) (eu-west-2) only in the Europe Regions listed in the preceding table. For

example, this applies to requests from Europe (Frankfurt) (eu-central-1) or Europe (Ireland) (eu-west-1).

- Amazon Quick never processes those requests in Europe (London) (eu-west-2) or Europe (Zurich) (eu-central-2).
- Amazon Quick can process a request that originates in Europe (London) (eu-west-2) in the listed Europe Regions or in Europe (London) itself.

 Note

Amazon Quick manages this cross-Region inference routing for you. You can't configure it, and Amazon Bedrock controls such as service control policies (SCPs) and IAM inference-profile permissions aren't available in Amazon Quick.

AWS Regions for Amazon Q in Quick

Topics

- [Supported AWS Regions for Amazon Q in Quick](#)

Supported AWS Regions for Amazon Q in Quick

Amazon Q in Quick Generative BI features including scenarios are currently supported in the following AWS Regions:

Region	Scenarios
US East (N. Virginia) (us-east-1)	✓ Yes
US East (Ohio) (us-east-2)	✓ Yes
US West (Oregon) (us-west-2)	✓ Yes
Asia Pacific (Mumbai) (ap-south-1)	✓ Yes

Region	Scenarios
Asia Pacific (Seoul) (ap-north-east-2)	✓ Yes
Asia Pacific (Singapore) (ap-southeast-1)	✓ Yes
Asia Pacific (Tokyo) (ap-north-east-1)	✓ Yes
Asia Pacific (Sydney) (ap-southeast-2)	✓ Yes
Canada (Central) (ca-central-1)	Not Available
Europe (Frankfurt) (eu-central-1)	✓ Yes
Europe (Ireland) (eu-west-1)	✓ Yes
Europe (London) (eu-west-2)	✓ Yes
Europe (Paris) (eu-west-3)	✓ Yes
Europe (Stockholm) (eu-north-1)	Not Available
Europe (Zurich) (eu-central-2)	Not Available
South America (São Paulo) (sa-east-1)	Not Available

For the AWS Regions where Amazon Quick processes cross-Region inference requests, see [Supported Regions for cross-Region inference](#).

Supported browsers

Before you start working with Amazon Quick, use the following table to verify that your browser is supported for Quick access.

Note

Amazon Quick ended support for Microsoft Internet Explorer 11 on July 31, 2022. We can no longer ensure that the features and web pages of Quick will function properly on IE 11. Please use one of our supported browsers: Microsoft Edge (Chromium), Google Chrome, or Mozilla Firefox.

Browser	Version	Check your version
Apple Safari	13 or later	Open Safari. On the menu, choose Safari , and then choose About Safari . The version number is shown in the dialog box that displays.
Google Chrome	Last three versions	Open Chrome and type chrome://version in your address bar. The version is in the Google Chrome field at the top of the results.
Microsoft Edge (Chromium)	Latest version	Not applicable.
Mozilla Firefox	Last three versions	Open Firefox. On the menu, choose the Help icon, and then choose About Firefox . The version number is listed underneath the Firefox name.

Setting up for Quick

This section helps you understand how to configure identity and access management, and control sign-up for users when you set up an Amazon Quick instance.

Note

For information about how to set up your Amazon Quick instance, including creating an AWS account, signing up for an Amazon Quick subscription, and signing in to Amazon Quick, see [Setting up and signing into Amazon Quick](#).

Topics

- [Using service control policies to restrict Amazon Quick sign-up options](#)
- [Identity and access management in Quick](#)
- [Allowlisting Amazon Quick domains](#)

Using service control policies to restrict Amazon Quick sign-up options

If you're an administrator in AWS Organizations, you can use service control policies (SCPs) to restrict how individuals in your organization can sign up for Amazon Quick. You can restrict the edition of Quick they can sign up for, and also the type of user that they can sign up for.

AWS Organizations is a user account management service that you can use to consolidate multiple AWS accounts into an organization that you create and centrally manage. You can use SCPs in AWS Organizations to manage the permissions in your organization. For more information, see [What is AWS Organizations?](#) and [Service control policies](#) in the *AWS Organizations User Guide*.

In the following topic, you can learn about two ways to restrict Quick sign-up options using SCPs in AWS Organizations. The topic includes an example SCP. To learn more about creating SCPs, see the following topics in the *AWS Organizations User Guide*:

- [Creating, updating, and deleting service control policies](#)
- [SCP syntax](#)
- [Strategies for using SCPs](#)

Topics

- [Restricting the Quick edition](#)
- [Restricting user management options](#)
- [Example SCP](#)
- [Restricting which accounts can create Quick subscriptions](#)

Restricting the Quick edition

To restrict the edition of Quick that your managed accounts can sign up for, use the `quicksight:Edition` condition key in your SCP. The values for this key are listed and described in the following table.

Key Name	Key Value	Description
<code>quicksight:Edition</code>	<code>standard</code>	Amazon Quick Standard Edition
	<code>enterprise</code>	Amazon Quick Enterprise Edition

Restricting user management options

To restrict the user management options that individuals in your organization can use to sign up for Quick, use the `quicksight:DirectoryType` condition key in your SCP. The values for this key are listed and described in the following table.

Key Name	Key Value	Description
<code>quicksight:DirectoryType</code>	<code>quicksight</code>	IAM federated identities and Amazon Quick-managed users
	<code>iam</code>	Only IAM federated identities
	<code>microsoft_ad</code>	Users managed in Microsoft Active Directory on AWS Directory Service for Microsoft Active Directory

Key Name	Key Value	Description
	ad_connector	Users managed in on-premises Active Directory and connected through AD_Connector to AWS Directory Service for Microsoft Active Directory
	iam_identity_center	Users managed in a Amazon Quick account that is integrated with IAM Identity Center.

Example SCP

The following example for Quick shows a service control policy that denies signing up for a Amazon Quick Standard Edition and prevents the ability to sign up using IAM Identity Center authentication. This policy uses the `quicksight:Subscribe` action, in addition to the condition keys previously described. For a list of Amazon Quick-specific keys for use in IAM permission policies, see [Actions, resources, and condition keys for Quick](#) in the *Service Authorization Reference*.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "Statement1",
      "Effect": "Deny",
      "Action": [
        "quicksight:Subscribe"
      ],
      "Resource": [
        "*"
      ],
      "Condition": {
        "ForAnyValue:StringEquals": {
          "quicksight:DirectoryType": [
            "iam_identity_center"
          ]
        }
      }
    }
  ]
}
```

```

    }
  },
  {
    "Sid": "Statement2",
    "Effect": "Deny",
    "Action": [
      "quicksight:Subscribe"
    ],
    "Resource": [
      "*"
    ],
    "Condition": {
      "StringEquals": {
        "quicksight:Edition": "standard"
      }
    }
  }
]
}

```

With this policy in effect, individuals in an organization can sign up only for Amazon Quick Enterprise Edition, and they must use authentication methods other than IAM Identity Center. If they try to sign up for Amazon Quick Standard Edition or attempt to use IAM Identity Center authentication, they will be restricted from signing up and receive a message explaining that they don't have the right permissions.

Restricting which accounts can create Quick subscriptions

In addition to restricting the edition and the user management options that individuals can sign up for, you can use an SCP to control *which* AWS accounts in your organization are allowed to create a Amazon Quick subscription. This helps you centralize Quick on a set of approved accounts and prevent individuals from starting new subscriptions in other accounts.

Two actions can create a Amazon Quick subscription: `quicksight:Subscribe` and `quicksight:CreateAccountSubscription`. If a policy denies either action, the deny takes effect and the attempt to create a subscription fails. To restrict subscription creation, include both actions in your Deny statement.

The following example SCP denies both actions in every account except the approved accounts that you list. The policy uses the `aws:PrincipalAccount` global condition key to compare the

account of the principal making the request against your approved accounts. Replace the example account IDs with the IDs of the accounts where you allow Quick subscriptions.

```
{
  "Version": "2012-10-17"      ,
  "Statement": [
    {
      "Sid": "Statement1",
      "Effect": "Deny",
      "Action": [
        "quicksight:Subscribe",
        "quicksight:CreateAccountSubscription"
      ],
      "Resource": [
        "*"
      ],
      "Condition": {
        "StringNotEquals": {
          "aws:PrincipalAccount": [
            "111111111111",
            "222222222222"
          ]
        }
      }
    }
  ]
}
```

With this policy in effect, only principals in the listed accounts can create a Amazon Quick subscription. Principals in any other account in your organization are prevented from signing up, and they receive a message explaining that they don't have the right permissions.

To allow subscription creation only through a specific provisioning identity – for example, a role that your infrastructure-as-code pipeline uses – use the `aws:PrincipalArn` global condition key instead of `aws:PrincipalAccount`. The following example denies all three actions unless the caller is the approved provisioning role. Replace the example role ARN with the ARN of your provisioning role.

```
{
  "Version": "2012-10-17"      ,
  "Statement": [
    {
```



```

    "Sid": "Statement1",
    "Effect": "Deny",
    "Action": [
        "quicksight:Subscribe",
        "quicksight:CreateAccountSubscription",
        "quicksight:CreateAdmin"
    ],
    "Resource": [
        "*"
    ],
    "Condition": {
        "ArnNotLike": {
            "aws:PrincipalArn": [
                "arn:aws:iam::*:role/QuickAccountProvisioner"
            ]
        }
    }
}
]
}

```

With this policy in effect, only the approved provisioning role can create a Amazon Quick subscription, regardless of which account the request comes from. Because CloudFormation creates resources as the principal that runs the template, you can use this approach to allow subscription creation only through CloudFormation. Run your template with the approved provisioning role. Protect the provisioning role by controlling who can assume it and who can change its trust policy.

This example also includes `quicksight:CreateAdmin`, which covers the Amazon Quick Standard Edition sign-up path in addition to `quicksight:Subscribe` and `quicksight:CreateAccountSubscription`.

As an alternative to using the condition key, you can attach a Deny statement for the same two actions (without the `aws:PrincipalAccount` condition) to every organizational unit (OU) except the OU that contains your approved accounts. For more information about where to attach policies, see [Strategies for using SCPs](#) in the *AWS Organizations User Guide*.

Identity and access management in Quick

Applies to: Enterprise Edition and Standard Edition

Intended audience: System administrators and Amazon Quick administrators

The following topics describe how to set up identity and access for Quick.

- [Using IAM Identity Center](#)
- [IAM federation](#)
- [Using Active Directory with Amazon Quick Enterprise edition](#)
- [Setting up IdP federation using IAM and Amazon Quick](#)
- [Using multi-factor authentication \(MFA\) with Amazon Quick](#)

 Note

In the following AWS Regions, Amazon Quick accounts can only use [IAM Identity Center](#) for identity and access management:

- af-south-1 Africa (Cape Town)
- ap-southeast-3 Asia Pacific (Jakarta)
- ap-southeast-5 Asia Pacific (Malaysia)
- eu-south-1 Europe (Milan)
- eu-south-2 Europe (Spain)
- eu-central-2 Europe (Zurich)
- il-central-1 Israel (Tel Aviv)
- me-central-1 Middle East (UAE)

The following sections help you configure the identity management method of your choice for Quick.

IAM permissions control access to some sections of the Amazon Quick administration console. The following table lists admin actions and whether they require IAM permissions.

Admin action	IAM permissions required
Account settings	 Yes

Admin action	IAM permissions required
Manage assets	✔ Yes
Amazon Q	✔ Yes
Manage subscriptions	✘ No
SPICE capacity	✘ No
Index capacity	✔ Yes
Manage users (view)	✘ No
Manage users > Role groups	✔ Yes
Manage domains	✘ No
Mobile settings	✘ No
Manage IP/VPC restrictions	✔ Yes
Manage VPC connections	✔ Yes
Manage OAuth client applications	✔ Yes
KMS keys	✔ Yes
AWS resources	✔ Yes
Default access policy	✔ Yes
IAM policy assignments	✔ Yes
AWS actions	✔ Yes
Extension access	✔ Yes
Custom permissions	✔ Yes
Configure SageMaker	✔ Yes

Admin action	IAM permissions required
Brand customization	✔ Yes
Agent customization	✔ Yes
Email customization	✔ Yes
Quick Usage Metrics	✔ Yes

If you have the Amazon Quick admin role, you can perform actions that do not require IAM permissions. To perform actions that require IAM permissions, sign in to the AWS Management Console as an IAM principal with the appropriate `quicksight:*` permissions. You can also perform some admin actions programmatically through the Amazon Quick API. For a list of available API operations, see the [Amazon Quick API Reference](#).

Topics

- [Using IAM](#)
- [Using IAM Identity Center](#)
- [IAM federation](#)
- [Using Active Directory with Amazon Quick Enterprise edition](#)
- [Using multi-factor authentication \(MFA\) with Amazon Quick](#)

Using IAM

AWS Identity and Access Management (IAM) is an AWS service that helps an administrator securely control access to AWS resources. IAM administrators control who can be *authenticated* (signed in) and *authorized* (have permissions) to use Amazon Quick resources. IAM is an AWS service that you can use with no additional charge.

Topics

- [Introduction to IAM concepts](#)
- [Using Quick with IAM](#)
- [Passing IAM roles to Quick](#)
- [IAM policy examples for Quick](#)

- [Provisioning users for Amazon Quick](#)
- [Troubleshooting Quick identity and access](#)

Introduction to IAM concepts

AWS Identity and Access Management (IAM) is an AWS service that helps an administrator to more securely control access to AWS resources. Administrators control who can be *authenticated* (signed in) and *authorized* (have permissions) to use Amazon Quick resources. IAM is an AWS service that you can use with no additional charge.

IAM is used with Amazon Quick in several ways, including the following:

- If your company uses IAM for their identity management, people might have IAM user names and passwords that they use to sign in to Amazon Quick.
- If you want your Amazon Quick users to be automatically created at first sign-in, you can use IAM to create a policy for users who are preauthorized to use Amazon Quick.
- If you want to create specialized access for specific groups of Amazon Quick users or to specific resources, you can use IAM policies to accomplish this.

Topics

- [Audience](#)
- [Authenticating with identities](#)
- [Managing access using policies](#)

Audience

Use the following to help understand the context of the information provided in this section, and how it applies to your role. How you use AWS Identity and Access Management (IAM) differs depending on the work that you do in Amazon Quick.

Service user – In some cases, you might use Amazon Quick as an Author or Reader to interact with data, analyses, and dashboards, spaces, and agents through the Amazon Quick by using the browser interface. In these cases, this section provides only background information for you. You don't directly interact with the IAM service, except if you use IAM to sign into Amazon Quick.

Amazon Quick administrator – If you're in charge of Amazon Quick resources at your company, you probably have full access to Amazon Quick. It's your job to determine which Amazon Quick

features and resources your team members should access. If you have specialized requirements that you can't solve by using the Amazon Quick admin panel, then you can work with your administrator to create permissions policies for your Amazon Quick users. To learn more about IAM, read this page to understand the basic concepts of IAM. To learn more about how your company can use IAM with Amazon Quick, see [Using Amazon Quick with IAM](#).

Administrator – If you're a system administrator, you might want to learn details about how you can write policies to manage access to Amazon Quick. To view examples of Amazon Quick identity-based policies that you can use in IAM, see [IAM identity-based policies for Amazon Quick](#).

Authenticating with identities

Authentication is how you sign in to AWS using your identity credentials. You must be authenticated as the AWS account root user, an IAM user, or by assuming an IAM role.

You can sign in as a federated identity using credentials from an identity source like AWS IAM Identity Center (IAM Identity Center), single sign-on authentication, or Google/Facebook credentials. For more information about signing in, see [How to sign in to your AWS account](#) in the *AWS Sign-In User Guide*.

For programmatic access, AWS provides an SDK and CLI to cryptographically sign requests. For more information, see [AWS Signature Version 4 for API requests](#) in the *IAM User Guide*.

Topics

- [AWS account root user](#)
- [IAM users and groups](#)
- [IAM roles](#)

AWS account root user

When you create an AWS account, you begin with one sign-in identity called the AWS account *root user* that has complete access to all AWS services and resources. We strongly recommend that you don't use the root user for everyday tasks. For tasks that require root user credentials, see [Tasks that require root user credentials](#) in the *IAM User Guide*.

IAM users and groups

An [IAM user](#) is an identity with specific permissions for a single person or application. We recommend using temporary credentials instead of IAM users with long-term credentials. For more

information, see [Require human users to use federation with an identity provider to access AWS using temporary credentials](#) in the *IAM User Guide*.

An [IAM group](#) specifies a collection of IAM users and makes permissions easier to manage for large sets of users. For more information, see [Use cases for IAM users](#) in the *IAM User Guide*.

IAM roles

An [IAM role](#) is an identity with specific permissions that provides temporary credentials. You can assume a role by [switching from a user to an IAM role \(console\)](#) or by calling an AWS CLI or AWS API operation. For more information, see [Methods to assume a role](#) in the *IAM User Guide*.

IAM roles are useful for federated user access, temporary IAM user permissions, cross-account access, cross-service access, and applications running on Amazon EC2. For more information, see [Cross account resource access in IAM](#) in the *IAM User Guide*.

Managing access using policies

You control access in AWS by creating policies and attaching them to AWS identities or resources. A policy defines permissions when associated with an identity or resource. AWS evaluates these policies when a principal makes a request. Most policies are stored in AWS as JSON documents. For more information about JSON policy documents, see [Overview of JSON policies](#) in the *IAM User Guide*.

Using policies, administrators specify who has access to what by defining which **principal** can perform **actions** on what **resources**, and under what **conditions**.

By default, users and roles have no permissions. An IAM administrator creates IAM policies and adds them to roles, which users can then assume. IAM policies define permissions regardless of the method used to perform the operation.

Identity-based policies

Identity-based policies are JSON permissions policy documents that you attach to an identity (user, group, or role). These policies control what actions identities can perform, on which resources, and under what conditions. To learn how to create an identity-based policy, see [Define custom IAM permissions with customer managed policies](#) in the *IAM User Guide*.

Identity-based policies can be *inline policies* (embedded directly into a single identity) or *managed policies* (standalone policies attached to multiple identities). To learn how to choose between

managed and inline policies, see [Choose between managed policies and inline policies](#) in the *IAM User Guide*.

Resource-based policies

Resource-based policies are JSON policy documents that you attach to a resource. Examples include *IAM role trust policies* and *Amazon S3 bucket policies*. In services that support resource-based policies, service administrators can use them to control access to a specific resource. You must [specify a principal](#) in a resource-based policy.

Resource-based policies are inline policies that are located in that service. You can't use AWS managed policies from IAM in a resource-based policy.

Access control lists (ACLs)

Access control lists (ACLs) control which principals (account members, users, or roles) have permissions to access a resource. ACLs are similar to resource-based policies, although they do not use the JSON policy document format.

Amazon S3, AWS WAF, and Amazon VPC are examples of services that support ACLs. To learn more about ACLs, see [Access control list \(ACL\) overview](#) in the *Amazon Simple Storage Service Developer Guide*.

Other policy types

AWS supports additional policy types that can set the maximum permissions granted by more common policy types:

- **Permissions boundaries** – Set the maximum permissions that an identity-based policy can grant to an IAM entity. For more information, see [Permissions boundaries for IAM entities](#) in the *IAM User Guide*.
- **Service control policies (SCPs)** – Specify the maximum permissions for an organization or organizational unit in AWS Organizations. For more information, see [Service control policies](#) in the *AWS Organizations User Guide*.
- **Resource control policies (RCPs)** – Set the maximum available permissions for resources in your accounts. For more information, see [Resource control policies \(RCPs\)](#) in the *AWS Organizations User Guide*.
- **Session policies** – Advanced policies passed as a parameter when creating a temporary session for a role or federated user. For more information, see [Session policies](#) in the *IAM User Guide*.

Multiple policy types

When multiple types of policies apply to a request, the resulting permissions are more complicated to understand. To learn how AWS determines whether to allow a request when multiple policy types are involved, see [Policy evaluation logic](#) in the *IAM User Guide*.

Using Quick with IAM

Applies to: Enterprise Edition and Standard Edition

Intended audience: System administrators

Before you use IAM to manage access to Amazon Quick, you should understand what IAM features are available to use with Amazon Quick. To get a high-level view of how Amazon Quick and other AWS services work with IAM, see [AWS Services That Work with IAM](#) in the *IAM User Guide*.

Topics

- [Amazon Quick Policies \(identity-based\)](#)
- [Amazon Quick policies \(resource-based\)](#)
- [Authorization based on Amazon Quick tags](#)
- [Amazon Quick IAM roles](#)

Amazon Quick Policies (identity-based)

With IAM identity-based policies, you can specify allowed or denied actions and resources as well as the conditions under which actions are allowed or denied. Amazon Quick supports specific actions, resources, and condition keys. To learn about all of the elements that you use in a JSON policy, see [IAM JSON Policy Elements Reference](#) in the *IAM User Guide*.

You can use AWS root credentials or IAM user credentials to create an Amazon Quick account. AWS root and administrator credentials already have all of the required permissions for managing Amazon Quick access to AWS resources.

However, we recommend that you protect your root credentials, and instead use IAM user credentials. To do this, you can create a policy and attach it to the IAM user and roles that you plan

to use for Amazon Quick. The policy must include the appropriate statements for the Amazon Quick administrative tasks you need to perform, as described in the following sections.

Important

Be aware of the following when working with Quick and IAM policies:

- Avoid directly modifying a policy that was created by Quick. When you modify it yourself, Quick can't edit it. This inability can cause an issue with the policy. To fix this issue, delete the previously modified policy.
- If you get an error on permissions when you try to create an Amazon Quick account, see [Actions Defined by Amazon Quick](#) in the *IAM User Guide*.
- In some cases, you might have an Amazon Quick account that you can't access even from the root account (for example, if you accidentally deleted its directory service). In this case, you can delete your old Amazon Quick account, then recreate it. For more information, see [Deleting your Amazon Quick subscription and closing the account](#).

Topics

- [Actions](#)
- [Resources](#)
- [Condition keys](#)
- [Examples](#)

Actions

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The Action element of a JSON policy describes the actions that you can use to allow or deny access in a policy. Include actions in a policy to grant permissions to perform the associated operation.

Policy actions in Amazon Quick use the following prefix before the action: `quicksight:`. For example, to grant someone permission to run an Amazon EC2 instance with the Amazon EC2 `RunInstances` API operation, you include the `ec2:RunInstances` action in their policy. Policy

statements must include either an `Action` or `NotAction` element. Amazon Quick defines its own set of actions that describe tasks that you can perform with this service.

To specify multiple actions in a single statement, separate them with commas as follows:

```
"Action": [  
    "quicksight:action1",  
    "quicksight:action2"]
```

You can specify multiple actions using wildcards (*). For example, to specify all actions that begin with the word `Create`, include the following action:

```
"Action": "quicksight:Create*"
```

Amazon Quick provides a number of AWS Identity and Access Management (IAM) actions. All Amazon Quick actions are prefixed with `quicksight:`, such as `quicksight:Subscribe`. For information about using Amazon Quick actions in an IAM policy, see [IAM policy examples for Amazon Quick](#).

To see the most up-to-date list of Amazon Quick actions, see [Actions Defined by Amazon Quick](#) in the *IAM User Guide*.

Resources

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The Resource JSON policy element specifies the object or objects to which the action applies. As a best practice, specify a resource using its [Amazon Resource Name \(ARN\)](#). For actions that don't support resource-level permissions, use a wildcard (*) to indicate that the statement applies to all resources.

```
"Resource": "*"
```

Following is an example policy. It means that the caller with this policy attached, is able to invoke the `CreateGroupMembership` operation on any group, provided that the user name they are adding to the group is not `user1`.

```
{
```

```
"Effect": "Allow",
"Action": "quicksight:CreateGroupMembership",
"Resource": "arn:aws:quicksight:us-east-1:aws-account-id:group/default/*",
"Condition": {
    "StringNotEquals": {
        "quicksight:UserName": "user1"
    }
}
```

Some Amazon Quick actions, such as those for creating resources, cannot be performed on a specific resource. In those cases, you must use the wildcard (*).

```
"Resource": "*"
```

Some API actions involve multiple resources. To specify multiple resources in a single statement, separate the ARNs with commas.

```
"Resource": [
    "resource1",
    "resource2"
```

To see a list of Amazon Quick resource types and their Amazon Resource Names (ARNs), see [Resources Defined by Amazon Quick](#) in the *IAM User Guide*. To learn with which actions you can specify the ARN of each resource, see [Actions Defined by Amazon Quick](#).

Condition keys

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The **Condition** element specifies when statements execute based on defined criteria. You can create conditional expressions that use [condition operators](#), such as equals or less than, to match the condition in the policy with values in the request. To see all AWS global condition keys, see [AWS global condition context keys](#) in the *IAM User Guide*.

Amazon Quick does not provide any service-specific condition keys, but it does support using some global condition keys. To see all AWS global condition keys, see [AWS Global Condition Context Keys](#) in the *IAM User Guide*.

Examples

To view examples of Amazon Quick identity-based policies, see [Amazon Quick Policies \(identity-based\)](#).

Amazon Quick policies (resource-based)

Amazon Quick doesn't support resource-based policies. However, you can use the Amazon Quick console to configure access to other AWS resources in your AWS account.

Authorization based on Amazon Quick tags

Amazon Quick does not support tagging resources or controlling access based on tags.

Amazon Quick IAM roles

An [IAM role](#) is an entity within your AWS account that has specific permissions. You can use IAM roles to group permissions together to make it easier to manage user's access to Amazon Quick actions.

Amazon Quick doesn't support the following role features:

- Service-linked roles.
- Service roles.
- Temporary credentials (direct use): However, Amazon Quick uses temporary credentials to allow users to assume an IAM role to access embedded dashboards. For more information, see [Embedded analytics for Amazon Quick](#).

For more information on how Amazon Quick uses IAM roles, see [Using Amazon Quick with IAM](#) and [IAM policy examples for Amazon Quick](#).

Passing IAM roles to Quick

Applies to: Enterprise Edition

When your IAM users sign up for Quick, they can choose to use the Amazon Quick-managed role (this is the default role). Or they can pass an existing IAM role to Amazon Quick.

Use the sections below to pass existing IAM roles to Amazon Quick

Topics

- [Prerequisites](#)

- [Attaching additional policies](#)
- [Using existing IAM roles in Quick](#)

Prerequisites

For your users to pass IAM roles to Amazon Quick, your administrator needs to complete the following tasks:

- **Create an IAM role.** For more information about creating IAM roles, see [Creating IAM roles](#) in the *IAM User Guide*.
- **Attach a trust policy to your IAM role that allows Amazon Quick to assume the role.** Use the following example to create a trust policy for the role. The following example trust policy allows the Quick principal to assume the IAM role that it's attached to.

For more information about creating IAM trust policies and attaching them to roles, see [Modifying a Role \(Console\)](#) in the *IAM User Guide*.

```
{
  "Version": "2012-10-17"      ,
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": "quicksight.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

- **Assign the following IAM permissions to your administrator (IAM users or roles):**
 - `quicksight:UpdateResourcePermissions` – This grants IAM users who are Amazon Quick administrators the permission to update resource-level permissions in Amazon Quick. For more information about resource types defined by Amazon Quick, see [Actions, resources, and condition keys for Quick](#) in the *IAM User Guide*.
 - `iam:PassRole` – This grants users permission to pass roles to Amazon Quick. For more information, see [Granting a user permissions to pass a role to an AWS service](#) in the *IAM User Guide*.

- `iam:ListRoles` – (Optional) This grants users permission to see a list of existing roles in Amazon Quick. If this permission is not provided, they can use an ARN to use existing IAM roles.

Following is an example IAM permissions policy that allows managing resource-level permissions, listing IAM roles, and passing IAM roles in Quick.

```
{
  "Version": "2012-10-17"
  ,
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "iam:ListRoles",
      "Resource": "arn:aws:iam::account-id:role:*"
    },
    {
      "Effect": "Allow",
      "Action": "iam:PassRole",
      "Resource": "arn:aws:iam::account-id:role/path/role-name",
      "Condition": {
        "StringEquals": {
          "iam:PassedToService": [
            "quicksight.amazonaws.com"
          ]
        }
      }
    },
    {
      "Effect": "Allow",
      "Action": "quicksight:UpdateResourcePermissions",
      "Resource": "*"
    }
  ]
}
```

For more examples of IAM policies that you can use with Amazon Quick, see [IAM policy examples for Amazon Quick](#).

For more information about assigning permissions policies to users or user groups, see [Changing permissions for an IAM user](#) in the *IAM User Guide*.

Attaching additional policies

If you're using another AWS service, such as Amazon Athena or Amazon S3, you can create a permissions policy that grants Amazon Quick permission to perform specific actions. You can then attach the policy to the IAM roles that you later pass to Amazon Quick. The following are examples of how you can set up and attach additional permissions policies to your IAM roles.

For an example managed policy for Amazon Quick in Athena, see [AWSQuicksightAthenaAccess Managed Policy](#) in the *Amazon Athena User Guide*. IAM users can access this role in Amazon Quick using the following ARN: `arn:aws:iam::aws:policy/service-role/AWSQuicksightAthenaAccess`.

The following is an example of a permissions policy for Amazon Quick in Amazon S3. For more information about using IAM with Amazon S3, see [Identity and access management in Amazon S3](#) in the *Amazon S3 User Guide*.

For information on how to create cross-account access from Amazon Quick to an Amazon S3 bucket in another account, see [How do I set up cross-account access from Quick to an Amazon S3 bucket in another account?](#) in the AWS Knowledge Center.

```
{
  "Version": "2012-10-17"
,
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "s3:ListAllMyBuckets",
      "Resource": "arn:aws:s3:::*"
    },
    {
      "Action": [
        "s3:ListBucket"
      ],
      "Effect": "Allow",
      "Resource": [
        "arn:aws:s3:::aws-athena-query-results-us-west-2-123456789"
      ]
    },
    {
      "Action": [
        "s3:GetObject",
        "s3:GetObjectVersion"
      ],
```



```

        "Effect": "Allow",
        "Resource": [
            "arn:aws:s3:::aws-athena-query-results-us-west-2-123456789/*"
        ],
    },
    {
        "Action": [
            "s3:ListBucketMultipartUploads",
            "s3:GetBucketLocation"
        ],
        "Effect": "Allow",
        "Resource": [
            "arn:aws:s3:::aws-athena-query-results-us-west-2-123456789"
        ],
    },
    {
        "Effect": "Allow",
        "Action": [
            "s3:PutObject",
            "s3:AbortMultipartUpload",
            "s3:ListMultipartUploadParts"
        ],
        "Resource": [
            "arn:aws:s3:::aws-athena-query-results-us-west-2-123456789/*"
        ],
    }
]
}

```

Using existing IAM roles in Quick

If you're a Amazon Quick administrator and have permissions to update Amazon Quick resources and pass IAM roles, you can use existing IAM roles in Amazon Quick. To learn more about the prerequisites for passing IAM roles in Amazon Quick, see the [Prerequisites](#) outlined in the previous list.

Use the following procedure to learn how to pass IAM roles in Amazon Quick.

To use an existing IAM role in Amazon Quick

1. In Amazon Quick, choose your account name in the navigation bar at top right and choose **Manage QuickSight**.

2. On the **Manage Amazon Quick** page that opens, choose **Security & Permissions** in the menu at left.
3. In the **Security & Permissions** page that opens, under **Amazon Quick access to AWS services**, choose **Manage**.
4. For **IAM role**, choose **Use an existing role**, and then do one of the following:
 - Choose the role that you want to use from the list.
 - Or, if you don't see a list of existing IAM roles, you can enter the IAM ARN for the role in the following format: `arn:aws:iam::account-id:role/path/role-name`.
5. Choose **Save**.

IAM policy examples for Quick

This section provides examples of IAM policies that you can use with Quick.

IAM identity-based policies for Quick

This section shows examples of identity-based policies to use with Quick.

Topics

- [IAM identity-based policies for Amazon Quick IAM console administration](#)

IAM identity-based policies for Amazon Quick IAM console administration

The following example shows the IAM permissions needed for Amazon Quick IAM console administration actions.

```
{
  "Version": "2012-10-17"
  ,
  "Statement": [
    {
      "Sid": "Statement1",
      "Effect": "Allow",
      "Action": [
        "quicksight:*",
        "iam:ListAttachedRolePolicies",
        "iam:GetPolicy",
        "iam:CreatePolicyVersion",
```

```

        "iam:DeletePolicyVersion",
        "iam:GetPolicyVersion",
        "iam:ListPolicyVersions",
        "iam:DeleteRole",
        "iam:CreateRole",
        "iam:GetRole",
        "iam:ListRoles",
        "iam:CreatePolicy",
        "iam:ListEntitiesForPolicy",
        "iam:listPolicies",
        "s3:ListAllMyBuckets",
        "athena:ListDataCatalogs",
        "athena:GetDataCatalog"
    ],
    "Resource": [
        "*"
    ]
}
]
}

```

IAM identity-based policies for Quick: dashboards

The following example shows an IAM policy that allows dashboard sharing and embedding for specific dashboards.

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": "quicksight:RegisterUser",
      "Resource": "*",
      "Effect": "Allow"
    },
    {
      "Action": "quicksight:GetDashboardEmbedUrl",
      "Resource": "arn:aws:quicksight:us-
west-2:111122223333:dashboard/1a1ac2b2-3fc3-4b44-5e5d-c6db6778df89",
      "Effect": "Allow"
    }
  ]
}

```

IAM identity-based policies for Quick: namespaces

The following examples show IAM policies that allow a Amazon Quick administrator to create or delete namespaces.

Creating namespaces

```
{
  "Version": "2012-10-17"      ,
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ds:AuthorizeApplication",
        "ds:UnauthorizeApplication",
        "ds>DeleteDirectory",
        "ds>CreateIdentityPoolDirectory",
        "ds:DescribeDirectories",
        "quicksight:CreateNamespace"
      ],
      "Resource": "*"
    }
  ]
}
```

Deleting namespaces

```
{
  "Version": "2012-10-17"      ,
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ds:UnauthorizeApplication",
        "ds>DeleteDirectory",
        "ds:DescribeDirectories",
        "quicksight>DeleteNamespace"
      ],
      "Resource": "*"
    }
  ]
}
```

IAM identity-based policies for Quick: custom permissions

The following example shows an IAM policy that allows a Amazon Quick administrator or a developer to manage custom permissions.

```
{
  "Version": "2012-10-17"      ,
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "quicksight:*CustomPermissions"
      ],
      "Resource": "*"
    }
  ]
}
```

The following example shows another way to grant the same permissions as shown in the previous example.

```
{
  "Version": "2012-10-17"      ,
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "quicksight:CreateCustomPermissions",
        "quicksight:DescribeCustomPermissions",
        "quicksight:ListCustomPermissions",
        "quicksight:UpdateCustomPermissions",
        "quicksight>DeleteCustomPermissions"
      ],
      "Resource": "*"
    }
  ]
}
```

IAM identity-based policies for Quick: customizing email report templates

The following example shows a policy that allows viewing, updating, and creating email report templates in Amazon Quick, as well as obtaining verification attributes for an Amazon Simple Email Service identity. This policy allows a Amazon Quick administrator to create and update custom email report templates, and to confirm that any custom email address they want to send email reports from is a verified identity in SES.

```
{
  "Version": "2012-10-17"
  ,
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "quicksight:DescribeAccountCustomization",
        "quicksight:CreateAccountCustomization",
        "quicksight:UpdateAccountCustomization",
        "quicksight:DescribeEmailCustomizationTemplate",
        "quicksight:CreateEmailCustomizationTemplate",
        "quicksight:UpdateEmailCustomizationTemplate",
        "ses:GetIdentityVerificationAttributes"
      ],
      "Resource": "*"
    }
  ]
}
```

IAM identity-based policies for Quick: create an Enterprise account with Amazon Quick managed users

The following example shows a policy that allows Amazon Quick admins to create an Enterprise edition Amazon Quick account with Amazon Quick managed users.

```
{
  "Version": "2012-10-17"
  ,
  "Statement": [
    {
      "Sid": "Statement1",
      "Effect": "Allow",
      "Action": [
        "quicksight:*",
        "iam:ListAttachedRolePolicies",

```

```

        "iam:GetPolicy",
        "iam:CreatePolicyVersion",
        "iam:DeletePolicyVersion",
        "iam:GetPolicyVersion",
        "iam:ListPolicyVersions",
        "iam:DeleteRole",
        "iam:CreateRole",
        "iam:GetRole",
        "iam:ListRoles",
        "iam:CreatePolicy",
        "iam:ListEntitiesForPolicy",
        "iam:listPolicies",
        "s3:ListAllMyBuckets",
        "athena:ListDataCatalogs",
        "athena:GetDataCatalog",
        "ds:AuthorizeApplication",
        "ds:UnauthorizeApplication",
        "ds:CheckAlias",
        "ds:CreateAlias",
        "ds:DescribeDirectories",
        "ds:DescribeTrusts",
        "ds:DeleteDirectory",
        "ds:CreateIdentityPoolDirectory"
    ],
    "Resource": [
        "*"
    ]
}
]
}

```

IAM identity-based policies for Quick: creating users

The following example shows a policy that allows creating Amazon Quick users only. For `quicksight:CreateReader`, `quicksight:CreateUser`, and `quicksight:CreateAdmin`, you can limit the permissions to **"Resource":**

"arn:aws:quicksight::<YOUR_AWS_ACCOUNTID>:user/\${aws:user_id}". For all other permissions described in this guide, use **"Resource": "*"**. The resource you specify limits the scope of the permissions to the specified resource.

```

{
    "Version": "2012-10-17"
    ,

```

```

    "Statement": [
      {
        "Action": [
          "quicksight:CreateUser"
        ],
        "Effect": "Allow",
        "Resource": "arn:aws:quicksight::<YOUR_AWS_ACCOUNTID>:user/${aws:userid}"
      }
    ]
  }
}

```

IAM identity-based policies for Quick: creating and managing groups

The following example shows a policy that allows Amazon Quick administrators and developers to create and manage groups.

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "quicksight:ListGroup",
        "quicksight:CreateGroup",
        "quicksight:SearchGroups",
        "quicksight:ListGroupMemberships",
        "quicksight:CreateGroupMembership",
        "quicksight>DeleteGroupMembership",
        "quicksight:DescribeGroupMembership",
        "quicksight:ListUsers"
      ],
      "Resource": "*"
    }
  ]
}

```

IAM identity-based policies for Quick: All access for Standard edition

The following example for Amazon Quick Standard edition shows a policy that allows subscribing and creating authors and readers. This example explicitly denies permission to unsubscribe from Amazon Quick.

```

{

```



```

"Version": "2012-10-17"
,
"Statement": [
  {
    "Effect": "Allow",
    "Action": [
      "ds:AuthorizeApplication",
      "ds:UnauthorizeApplication",
      "ds:CheckAlias",
      "ds:CreateAlias",
      "ds:DescribeDirectories",
      "ds:DescribeTrusts",
      "ds>DeleteDirectory",
      "ds>CreateIdentityPoolDirectory",
      "iam:ListAccountAliases",
      "quicksight:CreateUser",
      "quicksight:DescribeAccountSubscription",
      "quicksight:Subscribe"
    ],
    "Resource": "*"
  },
  {
    "Effect": "Deny",
    "Action": [
      "quicksight:Unsubscribe",
      "quicksight>DeleteAccountSubscription"
    ],
    "Resource": "*"
  }
]
}

```

IAM identity-based policies for Quick: All access for Enterprise edition with IAM Identity Center (Pro roles)

The following example for Amazon Quick Enterprise edition shows a policy that allows a Amazon Quick user to subscribe to Amazon Quick, create users, and manage Active Directory in a Amazon Quick account that is integrated with IAM Identity Center.

This policy also allows users to subscribe to Amazon Quick Pro roles that grant access to Amazon Q in Quick Generative BI capabilities. For more information about Pro roles in Amazon Quick, see [Get started with Generative BI](#).

This example explicitly denies permission to unsubscribe from Amazon Quick.

```

{
  "Statement": [
    {
      "Sid": "Statement1",
      "Effect": "Allow",
      "Action": [
        "quicksight:*",
        "iam:ListAttachedRolePolicies",
        "iam:GetPolicy",
        "iam:CreatePolicyVersion",
        "iam:DeletePolicyVersion",
        "iam:GetPolicyVersion",
        "iam:ListPolicyVersions",
        "iam:DeleteRole",
        "iam:CreateRole",
        "iam:GetRole",
        "iam:ListRoles",
        "iam:CreatePolicy",
        "iam:ListEntitiesForPolicy",
        "iam:listPolicies",
        "iam:CreateServiceLinkedRole",
        "s3:ListAllMyBuckets",
        "athena:ListDataCatalogs",
        "athena:GetDataCatalog",
        "sso:DescribeApplication",
        "sso:DescribeInstance",
        "sso:CreateApplication",
        "sso:PutApplicationAuthenticationMethod",
        "sso:PutApplicationGrant",
        "sso:DeleteApplication",
        "sso:SearchGroups",
        "sso:GetProfile",
        "sso:CreateApplicationAssignment",
        "sso:DeleteApplicationAssignment",
        "sso:ListInstances",
        "sso:DescribeRegisteredRegions",
        "organizations:DescribeOrganization",
        "user-subscriptions:CreateClaim",
        "user-subscriptions:UpdateClaim",
        "sso-directory:DescribeUser",
        "sso:ListApplicationAssignments",
        "sso-directory:DescribeGroup",
        "organizations:ListAWSServiceAccessForOrganization",

```

```

        "identitystore:DescribeUser",
        "identitystore:DescribeGroup"
    ],
    "Resource": [
        "*"
    ]
}
]
}

```

IAM identity-based policies for Quick: All access for Enterprise edition with IAM Identity Center

The following example for Amazon Quick Enterprise edition shows a policy that allows subscribing, creating users, and managing Active Directory in a Amazon Quick account that is integrated with IAM Identity Center.

This policy does not grant permissions to create Pro roles in Amazon Quick. To create a policy that grants permission to subscribe to Pro roles in Amazon Quick, see [IAM identity-based policies for Amazon Quick: All access for Enterprise edition with IAM Identity Center \(Pro roles\)](#).

This example explicitly denies permission to unsubscribe from Amazon Quick.

```

{
  "Statement": [
    {
      "Sid": "Statement1",
      "Effect": "Allow",
      "Action": [
        "quicksight:*",
        "iam:ListAttachedRolePolicies",
        "iam:GetPolicy",
        "iam:CreatePolicyVersion",
        "iam:DeletePolicyVersion",
        "iam:GetPolicyVersion",
        "iam:ListPolicyVersions",
        "iam>DeleteRole",
        "iam>CreateRole",
        "iam:GetRole",
        "iam:ListRoles",
        "iam:CreatePolicy",
        "iam:ListEntitiesForPolicy",
        "iam:listPolicies",
        "s3:ListAllMyBuckets",

```

```

        "athena:ListDataCatalogs",
        "athena:GetDataCatalog",
        "sso:DescribeApplication",
        "sso:DescribeInstance",
        "sso:CreateApplication",
        "sso:PutApplicationAuthenticationMethod",
        "sso:PutApplicationGrant",
        "sso:DeleteApplication",
        "sso:SearchGroups",
        "sso:GetProfile",
        "sso:CreateApplicationAssignment",
        "sso:DeleteApplicationAssignment",
        "sso:ListInstances",
        "sso:DescribeRegisteredRegions",
        "organizations:DescribeOrganization"
    ],
    "Resource": [
        "*"
    ]
}
]
}

```

IAM identity-based policies for Quick: all access for Enterprise edition with Active Directory

The following example for Amazon Quick Enterprise edition shows a policy that allows subscribing, creating users, and managing Active Directory in a Amazon Quick account that uses Active Directory for identity management. This example explicitly denies permission to unsubscribe from Amazon Quick.

```

{
  "Version": "2012-10-17"      ,
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ds:AuthorizeApplication",
        "ds:UnauthorizeApplication",
        "ds:CheckAlias",
        "ds:CreateAlias",
        "ds:DescribeDirectories",
        "ds:DescribeTrusts",
        "ds>DeleteDirectory",

```

```

        "ds:CreateIdentityPoolDirectory",
        "iam:ListAccountAliases",
        "quicksight:CreateAdmin",
        "quicksight:Subscribe",
        "quicksight:GetGroupMapping",
        "quicksight:SearchDirectoryGroups",
        "quicksight:SetGroupMapping"
    ],
    "Resource": "*"
  },
  {
    "Effect": "Deny",
    "Action": [
      "quicksight:Unsubscribe",
      "quicksight>DeleteAccountSubscription"
    ],
    "Resource": "*"
  }
]
}

```

IAM identity-based policies for Quick: active directory groups

The following example shows an IAM policy that allows Active Directory group management for an Amazon Quick Enterprise edition account.

```

{
  "Statement": [
    {
      "Action": [
        "ds:DescribeTrusts",
        "quicksight:GetGroupMapping",
        "quicksight:SearchDirectoryGroups",
        "quicksight:SetGroupMapping"
      ],
      "Effect": "Allow",
      "Resource": "*"
    }
  ],
  "Version": "2012-10-17"
}

```

IAM identity-based policies for Quick: using the admin asset management console

The following example shows an IAM policy that allows access to the admin asset management console.

```
{
  "Version": "2012-10-17"      ,
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "quicksight:SearchGroups",
        "quicksight:SearchUsers",
        "quicksight:ListNamespaces",
        "quicksight:DescribeAnalysisPermissions",
        "quicksight:DescribeDashboardPermissions",
        "quicksight:DescribeDataSetPermissions",
        "quicksight:DescribeDataSourcePermissions",
        "quicksight:DescribeFolderPermissions",
        "quicksight:ListAnalyses",
        "quicksight:ListDashboards",
        "quicksight:ListDataSets",
        "quicksight:ListDataSources",
        "quicksight:ListFolders",
        "quicksight:SearchAnalyses",
        "quicksight:SearchDashboards",
        "quicksight:SearchFolders",
        "quicksight:SearchDatasets",
        "quicksight:SearchDatasources",
        "quicksight:UpdateAnalysisPermissions",
        "quicksight:UpdateDashboardPermissions",
        "quicksight:UpdateDataSetPermissions",
        "quicksight:UpdateDataSourcePermissions",
        "quicksight:UpdateFolderPermissions"
      ],
      "Resource": "*"
    }
  ]
}
```

IAM identity-based policies for Quick: using the admin key management console

The following example shows an IAM policy that allows access to the admin key management console.

```
{
  "Version": "2012-10-17"
  ,
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "quicksight:DescribeKeyRegistration",
        "quicksight:UpdateKeyRegistration",
        "quicksight:ListKMSKeysForUser",
        "kms:CreateGrant",
        "kms:ListGrants",
        "kms:ListAliases"
      ],
      "Resource": "*"
    }
  ]
}
```

The "quicksight:ListKMSKeysForUser" and "kms:ListAliases" permissions are required to access customer managed keys from the Amazon Quick console.

"quicksight:ListKMSKeysForUser" and "kms:ListAliases" are not required to use the Amazon Quick key management APIs.

To specify which keys you want a user to be able to access, add the ARNs of the keys that you want the user to access to the UpdateKeyRegistration condition with the quicksight:KmsKeyArns condition key. Users can only access the keys specified in UpdateKeyRegistration. For more information about supported condition keys for Amazon Quick, see [Condition keys for Amazon Quick](#).

The example below grants Describe permissions for all CMKs that are registered to a Amazon Quick account and Update permissions to specific CMKs that are registered to the Amazon Quick account.

```
{
  "Version": "2012-10-17"
  ,
  "Statement": [
```

```

{
  "Effect": "Allow",
  "Action": [
    "quicksight:DescribeKeyRegistration"
  ],
  "Resource": "arn:aws:quicksight:us-west-2:123456789012:*"
},
{
  "Effect": "Allow",
  "Action": [
    "quicksight:UpdateKeyRegistration"
  ],
  "Resource": "arn:aws:quicksight:us-west-2:123456789012:*",
  "Condition": {
    "ForAllValues:StringEquals": {
      "quicksight:KmsKeyArns": [
        "arn:aws:kms:us-west-2:123456789012:key/key-id-of-key1",
        "arn:aws:kms:us-west-2:123456789012:key/key-id-of-key2",
        "..."
      ]
    }
  }
},
{
  "Effect": "Allow",
  "Action": [
    "kms:CreateGrant",
    "kms:ListGrants"
  ],
  "Resource": "arn:aws:kms:us-west-2:123456789012:key/*"
}
]
}

```

AWS resources Quick: scoping policies in Enterprise edition

The following example for Amazon Quick Enterprise edition shows a policy that allows setting default access to AWS resources and scoping policies for permissions to AWS resources.

```

{
  "Version": "2012-10-17"
  ,
  "Statement": [
    {

```



```
        "Action": [
            "quicksight:*IAMPolicyAssignment*",
            "quicksight:AccountConfigurations"
        ],
        "Effect": "Allow",
        "Resource": "*"
    }
}
```

Provisioning users for Amazon Quick

Applies to: Enterprise Edition and Standard Edition

Intended audience: System administrators and Amazon Quick administrators

Self-provisioning an Amazon Quick administrator

Amazon Quick administrators are users who can also manage Amazon Quick features such as account settings and accounts. They can also purchase additional Amazon Quick user subscriptions, purchase [SPICE](#), and cancel the subscription to Amazon Quick for your AWS account.

You can use an AWS user or group policy to give users the ability to add themselves as administrators of Amazon Quick. Users that have been granted this ability can only add themselves as administrators and can't use this policy to add others. Their accounts become active and billable the first time that they open Amazon Quick. To set up self-provisioning, give these users permission to use the `quicksight:CreateAdmin` action.

Alternatively, you can use the following procedure to use the console to set or create the administrator for Amazon Quick.

To make a user the Amazon Quick administrator

1. Create the AWS user:
 - Use IAM to create the user that you want to be the administrator of Amazon Quick. Alternatively, identify an existing user in IAM for the administrator role. You can also put the user inside a new group, for manageability.

- Grant the user (or group) sufficient permissions.
2. Sign in to your AWS Management Console with the target user's credentials.
 3. Go to <http://quicksight.aws.amazon.com/sn/console/get-user-email>, type in the target user's email address, and choose **Continue**.

On success, the target user is now an administrator in Amazon Quick.

Self-provisioning an Amazon Quick author

Amazon Quick authors can create data sources, datasets, analyses, and dashboards. They can share analyses and dashboards with other Amazon Quick users in your Amazon Quick account. However, they don't have access to the **Manage Amazon Quick** menu. They can't change account settings, manage accounts, purchase additional Amazon Quick user subscriptions or [SPICE](#) capacity, or cancel the subscription to Amazon Quick for your AWS account. Author Pro users can additionally create content using natural language, build knowledge bases, configure actions, and access advanced automation capabilities.

You can use an AWS user or group policy to give users the ability to create an Amazon Quick author account for themselves. Their accounts become active and billable the first time they open Amazon Quick. To set up self-provisioning, you need to give them permission to use the `quicksight:CreateUser` action.

Self-provisioning an Amazon Quick read-only user

Amazon Quick read-only users or *readers* can view and manipulate dashboards that are shared with them, but they can't make any changes or save a dashboard for further analysis. Amazon Quick readers can't create data sources, datasets, analyses, or visuals. They can't do any administrative tasks. Choose this role for people who are consumers of the dashboards but don't author their own analysis, for example, executives. Reader Pro users have access to advanced features including AI chat agents, collaborative spaces, flows, and extensions.

If you are using Microsoft Active Directory with Amazon Quick, you can manage read-only permissions by using a group. Otherwise, you can bulk-invite users to use Amazon Quick. You can also use an AWS user or group policy to give people the ability to create an Amazon Quick reader account for themselves.

Reader accounts become active and billable the first time they open Amazon Quick. If you decide to upgrade or downgrade a user, billing for that user is prorated for the month. To set up self-

provisioning, you need to give them permission to use the `quicksight:CreateReader` action. For complete details about when billing begins, proration, and what happens when you remove a user, see [Understanding Amazon Quick user billing](#).

Readers that are used to automatically or programmatically refresh dashboards for near real-time use cases must choose capacity pricing. For readers under user pricing, each reader is limited to manual use by one individual only.

Troubleshooting Quick identity and access

Applies to: Enterprise Edition and Standard Edition

Intended audience: System administrators

Use the following information to help you diagnose and fix common issues that you might encounter when working with Amazon Quick and IAM.

Topics

- [I am not authorized to perform an action in Amazon Quick](#)
- [I am not authorized to perform iam:PassRole](#)
- [I want to allow people outside of my AWS account to access my Amazon Quick resources](#)

I am not authorized to perform an action in Amazon Quick

If the AWS Management Console tells you that you're not authorized to perform an action, then you must contact your administrator for assistance.

The following example error occurs when the `mateojackson` IAM user tries to use the console to view details about a `widget` but does not have `quicksight:GetWidget` permissions.

```
User: arn:aws:iam::123456789012:user/mateojackson is not authorized to perform:
quicksight:GetWidget on resource: my-example-widget
```

In this case, Mateo asks his administrator to update his policies to allow him to access the `my-example-widget` resource using the `quicksight:GetWidget` action.

I am not authorized to perform iam:PassRole

If you receive an error that you're not authorized to perform the `iam:PassRole` action, your policies must be updated to allow you to pass a role to Amazon Quick.

Some AWS services allow you to pass an existing role to that service instead of creating a new service role or service-linked role. To do this, you must have permissions to pass the role to the service.

The following example error occurs when an IAM user named `marymajor` tries to use the console to perform an action in Amazon Quick. However, the action requires the service to have permissions that are granted by a service role. Mary does not have permissions to pass the role to the service.

```
User: arn:aws:iam::123456789012:user/marymajor is not authorized to perform:
iam:PassRole
```

In this case, Mary's policies must be updated to allow her to perform the `iam:PassRole` action.

If you need help, contact your AWS administrator. Your administrator is the person who provided you with your sign-in credentials.

I want to allow people outside of my AWS account to access my Amazon Quick resources

You can create a role that users in other accounts or people outside of your organization can use to access your resources. You can specify who is trusted to assume the role. For services that support resource-based policies or access control lists (ACLs), you can use those policies to grant people access to your resources.

To learn more, consult the following:

- To learn whether Amazon Quick supports these features, see [Using Quick with IAM](#).
- To learn how to provide access to your resources across AWS accounts that you own, see [Providing access to an IAM user in another AWS account that you own](#) in the *IAM User Guide*.
- To learn how to provide access to your resources to third-party AWS accounts, see [Providing access to AWS accounts owned by third parties](#) in the *IAM User Guide*.
- To learn how to provide access through identity federation, see [Providing access to externally authenticated users \(identity federation\)](#) in the *IAM User Guide*.
- To learn the difference between using roles and resource-based policies for cross-account access, see [Cross account resource access in IAM](#) in the *IAM User Guide*.

Using IAM Identity Center

Applies to: Enterprise Edition and Standard Edition

Intended audience: System administrators and Amazon Quick administrators

IAM Identity Center helps you securely create or connect your workforce identities and manage their access across AWS accounts and applications.

Before you integrate your Amazon Quick account with IAM Identity Center, set up IAM Identity Center in your AWS account. If you haven't set up IAM Identity Center in your AWS organization, see [Getting started](#) in the *AWS IAM Identity Center User Guide*.

If you want to configure an external identity provider with IAM Identity Center, see [Supported identity providers](#) to view a list of supported identity providers' configuration steps.

Topics

- [Configure your Amazon Quick account with IAM Identity Center](#)

Configure your Amazon Quick account with IAM Identity Center

Applies to: Enterprise Edition

Intended audience: System administrators

IAM Identity Center helps you securely create or configure your existing workforce identities and manage their access across AWS accounts and applications. IAM Identity Center is the recommended approach for workforce authentication and authorization on AWS for organizations of any size and type. To learn more about IAM Identity Center, see [AWS IAM Identity Center](#).

Configure Amazon Quick and IAM Identity Center so that you can sign up for a new Amazon Quick account with an IAM Identity Center configured identity source. With IAM Identity Center, you can configure your external identity provider as an identity source. You can also use IAM Identity Center

as an identity store if you don't want to use a third-party identity provider with Amazon Quick. Identity methods can't be changed after your account is created.

When you integrate your Amazon Quick account with IAM Identity Center, Amazon Quick account administrators can create a new Amazon Quick account that automatically has the identity provider's groups available. This simplifies asset sharing at scale in Amazon Quick.

To learn more how to sign up for an Amazon Quick account with IAM Identity Center, see [Signing up for an Amazon Quick subscription](#).

Set up Amazon Quick with IAM Identity Center multi-Region

You can create an Amazon Quick Enterprise subscription in a supported AWS Region that is configured as an additional Region for your IAM Identity Center instance. When you call `CreateAccountSubscription` against a target Region endpoint, that Region becomes the capacity Region for your Quick default namespace.

This procedure applies only to the IAM Identity Center authentication method. It does not create a second Quick subscription or replicate an existing one. Each AWS account supports one Quick account subscription.

Region concepts

The following Region types are relevant to this procedure:

- **IAM Identity Center primary Region** – The Region where your IAM Identity Center instance was originally created.
- **IAM Identity Center additional Region** – A Region to which your IAM Identity Center instance is replicated and active.
- **Quick capacity Region** – The Region where your Quick default namespace is provisioned. This is determined by the endpoint Region where you call `CreateAccountSubscription`.

Prerequisites

For detailed instructions, see [Multi-Region IAM Identity Center](#) and [Replicate to an additional Region](#) in the *AWS IAM Identity Center User Guide*.

Before you create your Quick subscription in an additional Region, complete the following steps in the order shown:

1. Configure a multi-Region customer managed KMS key for your IAM Identity Center instance in the primary Region.
2. Create a replica of the customer managed KMS key in the target Region.
3. Add the target Region to IAM Identity Center and wait until the Region status is ACTIVE.
4. If your organization uses an external identity provider, add the target Region's Assertion Consumer Service (ACS) URL to your identity provider configuration. For more information, see [Multi-Region workforce access](#) in the *AWS IAM Identity Center User Guide*.
5. Confirm that the target Region is a supported Quick signup Region. For a list of supported Regions, see [Supported AWS Regions for Amazon Quick](#).

To create a subscription in an additional Region

Use the AWS Command Line Interface or an AWS SDK to create an Amazon Quick Enterprise subscription in the target Region.

To create the subscription (AWS CLI)

1. Configure the AWS CLI to use the target Region. The configured Region determines the `CreateAccountSubscription` endpoint.
2. Run the following command. Replace the placeholder values with your own.

```
aws quicksight create-account-subscription \
  --aws-account-id <ACCOUNT_ID> \
  --account-name "<ACCOUNT_NAME>" \
  --notification-email <EMAIL> \
  --edition ENTERPRISE \
  --authentication-method IAM_IDENTITY_CENTER \
  --iam-identity-center-instance-arn <IDC_INSTANCE_ARN> \
  --admin-group "<ADMIN_GROUP_NAME>" \
  --region <TARGET_REGION>
```

The following list describes the required and recommended parameters:

- `--aws-account-id` – Your 12-digit AWS account ID.
- `--account-name` – A name for the Quick account.
- `--notification-email` – The email address for service notifications.
- `--edition ENTERPRISE` – IAM Identity Center authentication requires Enterprise edition.

- `--authentication-method IAM_IDENTITY_CENTER` – Specifies IAM Identity Center as the authentication method.
- `--iam-identity-center-instance-arn` – The ARN of your IAM Identity Center instance.
- `--admin-group` or `--admin-pro-group` – The name of the IAM Identity Center group to assign as Quick administrators.
- `--region` – The target Region for the Quick capacity Region. This must be an active IAM Identity Center additional Region and a supported Quick signup Region.

Verify the subscription

After the command completes, confirm that the subscription exists and is active.

To verify the subscription (AWS CLI)

1. Run the following command to check the account status:

```
aws quicksight describe-account-subscription \  
  --aws-account-id <ACCOUNT_ID> \  
  --region <TARGET_REGION>
```

Confirm that `AccountSubscriptionStatus` is `ACCOUNT_CREATED`.

2. Run the following command to verify the namespace and capacity Region:

```
aws quicksight list-namespaces \  
  --aws-account-id <ACCOUNT_ID> \  
  --region <TARGET_REGION>
```

Confirm that `CreationStatus` is `CREATED` and that `CapacityRegion` matches the target Region.

Considerations

Review the following considerations:

- You must use Amazon Quick Enterprise edition for IAM Identity Center authentication.

- You must specify either `--admin-group` or `--admin-pro-group` when you use IAM Identity Center authentication.
- You must choose a target Region that is both an active IAM Identity Center additional Region and a supported Quick signup Region.
- You can create only one Quick account subscription for each AWS account. If your account already has a subscription, `CreateAccountSubscription` returns a `resource-exists` error.
- You cannot change the authentication method after the Quick account is created.

Considerations

Irreversible actions

The following actions permanently prevent all users from signing in to your Amazon Quick account. You cannot undo these actions.

- Disabling or deleting the Amazon Quick application in the IAM Identity Center console. If you want to delete your Amazon Quick account, see [Closing your Amazon Quick account](#).
- Migrating the Amazon Quick account that contains your IAM Identity Center configuration to an AWS Organization that does not contain the IAM Identity Center instance that your Amazon Quick account is configured to.
- Deleting the IAM Identity Center instance that is configured to your Amazon Quick account.
- Editing IAM Identity Center application attributes, for example the **requires assignment** attribute.

IAM federation

Applies to: Enterprise Edition and Standard Edition

Intended audience: System administrators

⚠ Important

Amazon Quick recommends that you integrate new Amazon Quick subscriptions with IAM Identity Center for identity management. This IAM identity federation user guide is provided as a reference for existing account configurations. For more information on integrating your Amazon Quick account with IAM Identity Center, see [Configure your Amazon Quick account with IAM Identity Center](#).

ℹ Note

IAM identity federation doesn't support syncing identity provider groups with Amazon Quick.

Amazon Quick supports identity federation in both Standard and Enterprise editions. When you use federated users, you can manage users with your enterprise identity provider (IdP) and use AWS Identity and Access Management (IAM) to authenticate users when they sign in to Quick. You can use a third-party identity provider that supports Security Assertion Markup Language 2.0 (SAML 2.0) to provide an onboarding flow for your Amazon Quick users. Such identity providers include Microsoft Active Directory Federation Services, Okta, and Ping One Federation Server. With identity federation, your users get one-click access to their Amazon Quick applications using their existing identity credentials. You also have the security benefit of identity authentication by your identity provider. You can control which users have access to Amazon Quick using your existing identity provider.

Topics

- [Initiating sign-on from the identity provider \(IdP\)](#)
- [Setting up IdP federation using IAM and Amazon Quick](#)
- [Initiating sign-on from Quick](#)
- [Setting up service provider-initiated federation with Quick Enterprise edition](#)
- [Configuring email syncing for federated users in Quick](#)
- [Tutorial: Amazon Quick and IAM identity federation](#)

Initiating sign-on from the identity provider (IdP)

Applies to: Enterprise Edition and Standard Edition

Intended audience: System administrators

Note

IAM identity federation doesn't support syncing identity provider groups with Amazon Quick.


In this scenario, your users initiate the sign-on process from the identity provider's portal. After the users are authenticated, they sign in to Amazon Quick. After Quick checks that they are authorized, your users can access Quick.

Beginning with a user signing into the IdP, authentication flows through these steps:

1. The user browses to `https://applications.example.com` and signs on to the IdP. At this point, the user isn't signed in to the service provider.
2. The federation service and the IdP authenticate the user:
 - a. The federation service requests authentication from the organization's identity store.
 - b. The identity store authenticates the user and returns the authentication response to the federation service.
 - c. When authentication is successful, the federation service posts the SAML assertion to the user's browser.
3. The user opens Amazon Quick:
 - a. The user's browser posts the SAML assertion to the AWS Sign-In SAML endpoint (`https://signin.aws.amazon.com/saml`).
 - b. AWS Sign-In receives the SAML request, processes the request, authenticates the user, and forwards the authentication token to the Amazon Quick service.
4. Amazon Quick accepts the authentication token from AWS and presents Amazon Quick to the user.

From the user's perspective, the process happens transparently. The user starts at your organization's internal portal and lands at an Amazon Quick application portal, without ever having to supply any AWS credentials.

In the following diagram, you can find an authentication flow between Amazon Quick and a third-party identity provider (IdP). In this example, the administrator has set up a sign-in page to access Amazon Quick, called `applications.example.com`. When a user signs in, the sign-in page posts a request to a federation service that complies with SAML 2.0. The end user initiates authentication from the sign-on page of the IdP.



For information from some common providers, see the following third-party documentation:

- CA – [Enabling SAML 2.0 HTTP Post Binding](#)
- Okta – [Planning a SAML deployment](#)
- Ping – [Amazon integrations](#)

Use the following topics to understand using an existing federation with AWS:

- [Identity federation in AWS](#) on the AWS website

- [Providing access to externally authenticated users \(identity federation\)](#) in the *IAM User Guide*
- [Enabling SAML 2.0 federated users to access the AWS Management Console](#) in the *IAM User Guide*

Setting up IdP federation using IAM and Amazon Quick

Applies to: Enterprise Edition and Standard Edition

Intended audience: System administrators

Note

IAM identity federation doesn't support syncing identity provider groups with Amazon Quick.

You can use an AWS Identity and Access Management (IAM) role and a relay state URL to configure an identity provider (IdP) that is compliant with SAML 2.0. The role grants users permissions to access Amazon Quick. The relay state is the portal that the user is forwarded to, after successful authentication by AWS.

Topics

- [Prerequisites](#)
- [Step 1: Create a SAML provider in AWS](#)
- [Step 2: Configure permissions in AWS for your federated users](#)
- [Step 3: Configure the SAML IdP](#)
- [Step 4: Create assertions for the SAML authentication response](#)
- [Step 5: Configure the relay state of your federation](#)

Prerequisites

Before configuring your SAML 2.0 connection, do the following:

- Configure your IdP to establish a trust relationship with AWS:

- Inside your organization's network, configure your identity store, such as Windows Active Directory, to work with a SAML-based IdP. SAML-based IdPs include Active Directory Federation Services, Shibboleth, and so on.
- Using your IdP, generate a metadata document that describes your organization as an identity provider.
- Set up SAML 2.0 authentication, using the same steps as for the AWS Management Console. When this process is complete, you can configure your relay state to match the relay state of Quick. For more information, see [Configure the relay state of your federation](#).
- Create an Amazon Quick account and note the name to use when you configure your IAM policy and IdP. For more information on creating an Amazon Quick account, see [Signing up for an Amazon Quick subscription](#).

After you create the setup to federate to the AWS Management Console as outlined in the tutorial, you can edit the relay state provided in the tutorial. You do so with the relay state of Amazon Quick, described in step 5 following.

For more information, see the following resources:

- [Integrating Third-Party SAML Solution Providers with AWS](#) in the *IAM User Guide*.
- [Troubleshooting SAML 2.0 federation with AWS](#), also in the *IAM User Guide*.
- [Setting up trust between ADFS and AWS and using Active Directory credentials to connect to Amazon Athena with ODBC driver](#) – This walkthrough article is helpful, although you don't need to set up Athena in order to use Amazon Quick.

Step 1: Create a SAML provider in AWS

Your SAML identity provider defines your organization's IdP to AWS. It does so by using the metadata document that you previously generated using your IdP.

To create a SAML provider in AWS

1. Sign in to the AWS Management Console and open the IAM console at <https://console.aws.amazon.com/iam/>.
2. Create a new SAML provider, which is an entity in IAM that holds information about your organization's identity provider. For more information, see [Creating SAML Identity Providers](#) in the *IAM User Guide*.

3. As part of this process, upload the metadata document produced by the IdP software in your organization noted in the previous section.

Step 2: Configure permissions in AWS for your federated users

Next, create an IAM role that establishes a trust relationship between IAM and your organization's IdP. This role identifies your IdP as a principal (trusted entity) for the purposes of federation. The role also defines which users authenticated by your organization's IdP are allowed to access Amazon Quick. For more information about creating a role for a SAML IdP, see [Creating a Role for SAML 2.0 Federation](#) in the *IAM User Guide*.

After you have created the role, you can limit the role to have permissions only to Amazon Quick by attaching an inline policy to the role. The following sample policy document provides access to Amazon Quick. This policy allows the user access to Amazon Quick and allows them to create both author accounts and reader accounts.

Note

In the following example, replace `<YOUR_AWS_ACCOUNT_ID>` with your 12-digit AWS account ID (with no hyphens '-').

```
{
  "Statement": [
    {
      "Action": [
        "quicksight:CreateUser"
      ],
      "Effect": "Allow",
      "Resource": [
        "arn:aws:quicksight::<YOUR_AWS_ACCOUNT_ID>:user/${aws:userid}"
      ]
    }
  ],
  "Version": "2012-10-17"
}
```

If you want to provide access to Amazon Quick and also the ability to create Amazon Quick admins, authors (standard users), and readers, you can use the following policy example.

```
{
  "Statement": [
    {
      "Action": [
        "quicksight:CreateAdmin"
      ],
      "Effect": "Allow",
      "Resource": [
        "arn:aws:quicksight::<YOUR_AWS_ACCOUNT_ID>:user/${aws:userid}"
      ]
    }
  ],
  "Version": "2012-10-17"
}
```

You can view account details in the AWS Management Console.

After you have set up SAML and the IAM policy or policies, you don't need to invite users manually. The first time that users open Amazon Quick, they are provisioned automatically, using the highest level permissions in the policy. For example, if they have permissions to both `quicksight:CreateUser` and `quicksight:CreateReader`, they are provisioned as authors. If they also have permissions to `quicksight:CreateAdmin`, they are provisioned as admins. Each permission level includes the ability to create the same level user and below. For example, an author can add other authors or readers.

Users who are invited manually are created in the role assigned by the person who invited them. They don't need to have policies that grant them permissions.

Step 3: Configure the SAML IdP

After you create the IAM role, update your SAML IdP about AWS as a service provider. To do so, install the `saml-metadata.xml` file found at <https://signin.aws.amazon.com/static/saml-metadata.xml>.

To update the IdP metadata, see the instructions provided by your IdP. Some providers give you the option to type the URL, after which the IdP gets and installs the file for you. Others require you to download the file from the URL and then provide it as a local file.

For more information, see your IdP documentation.

Step 4: Create assertions for the SAML authentication response

Next, configure the information that the IdP passes as SAML attributes to AWS as part of the authentication response. For more information, see [Configuring SAML Assertions for the Authentication Response](#) in the *IAM User Guide*.

Step 5: Configure the relay state of your federation

Finally, configure the relay state of your federation to point to the Amazon Quick relay state URL. After successful authentication by AWS, the user is directed to Amazon Quick, defined as the relay state in the SAML authentication response.

The relay state URL for Amazon Quick is as follows.

```
https://quicksight.aws.amazon.com
```

Initiating sign-on from Quick

Applies to: Enterprise Edition

Intended audience: System administrators

Note

IAM identity federation doesn't support syncing identity provider groups with Amazon Quick.

In this scenario, your user initiates the sign-on process from an Amazon Quick application portal without being signed on to the identity provider. In this case, the user has a federated account managed by a third-party IdP. The user might have a user account on Quick. Quick sends an authentication request to the IdP. After the user is authenticated, Quick opens.

Beginning with the user signing into Quick, authentication flows through these steps:

1. The user opens Quick. At this point, the user isn't signed in to the IdP.
2. The user attempts to sign in to Amazon Quick.
3. Amazon Quick redirects the user's input to the federation service and requests authentication.
4. The federation service and the IdP authenticate the user:
 - a. The federation service requests authentication from the organization's identity store.
 - b. The identity store authenticates the user and returns the authentication response to the federation service.
 - c. When authentication is successful, the federation service posts the SAML assertion to the user's browser.
 - d. The user's browser posts the SAML assertion to the AWS Sign-In SAML endpoint (<https://signin.aws.amazon.com/saml>).
 - e. AWS Sign-In receives the SAML request, processes the request, authenticates the user, and forwards the authentication token to the Amazon Quick service.
5. Amazon Quick accepts the authentication token from AWS and presents Amazon Quick to the user.

From the user's perspective, the process happens transparently. The user starts at an Amazon Quick application portal. Amazon Quick negotiates authentication with your organization's federation service and AWS. Amazon Quick opens, without the user needing to supply any additional credentials.

Setting up service provider–initiated federation with Quick Enterprise edition

Applies to: Enterprise Edition

Intended audience: System administrators

Note

IAM identity federation doesn't support syncing identity provider groups with Amazon Quick.

After you have finished configuring your identity provider with AWS Identity and Access Management (IAM), you can set up service provider–initiated sign in through Amazon Quick Enterprise Edition. For Quick-initiated IAM federation to work, you need to authorize Quick to send the authentication request to your IdP. A Quick administrator can configure this by adding the following information provided by the IdP:

- The IdP URL – Quick redirects users to this URL for authentication.
- The relay state parameter – This parameter relays the state that the browser session was in when it was redirected for authentication. The IdP redirects the user back to the original state after authentication. The state is provided as a URL.

The following table shows the standard authentication URL and relay state parameter for redirecting the user to the Quick URL that you provide.

Identity provider	Parameter	Authentication URL
Auth0	RelayState	<code>https://<sub_domain>.auth0.com/samlp/<app_id></code>
Google accounts	RelayState	<code>https://accounts.google.com/o/saml2/initssso?idpid=<idp_id>&spid=<sp_id>&forceauthn=false</code>
Microsoft Azure	RelayState	<code>https://myapps.microsoft.com/signin/<app_name>/<app_id>?tenantId=<tenant_id></code>
Okta	RelayState	<code>https://<sub_domain>.okta.com/app/<app_name>/<app_id>/sso/saml</code>
PingFederate	TargetResource	<code>https://<host>/idp/<idp_id>/startSSO.ping?PartnerSpId=<sp_id></code>
PingOne	TargetResource	<code>https://sso.connect.pingidentity.com/sso/sp/initssso?saasid=<app_id>&idpid=<idp_id></code>

Amazon Quick supports connecting to one IdP per AWS account. The configuration page in Amazon Quick provides you with test URLs based on your entries, so you can test the settings

before you turn on the feature. To make the process even more seamless, Amazon Quick provides a parameter (`enable-sso=0`) to temporarily turn off Amazon Quick initiated IAM federation, in case you need to disable it temporarily.

To set up Amazon Quick as a service provider that can initiate IAM federation for an existing IdP

1. Make sure that you already have IAM federation set up in your IdP, in IAM, and Amazon Quick. To test this setup, check if you can share a dashboard with another person in your company's domain.
2. Open Amazon Quick, and choose **Manage Amazon Quick** from your profile menu at upper right.

To perform this procedure, you need to be a Amazon Quick administrator. If you aren't, you can't see **Manage Amazon Quick** under your profile menu.

3. Choose **Single sign-on (IAM federation)** from the navigation pane.
4. For **Configuration**, **IdP URL**, enter the URL that your IdP provides to authenticate users.
5. For **IdP URL**, enter the parameter that your IdP provides to relay state, for example `RelayState`. The actual name of the parameter is provided by your IdP.
6. Test signing in:
 - To test signing in with your identity provider, use the custom URL provided in **Test starting with your IdP**. You should arrive at the start page for Amazon Quick, for example `https://quicksight.aws.amazon.com/sn/start`.
 - To test signing in with Amazon Quick first, use the custom URL provided in **Test the end-to-end experience**. The `enable-sso` parameter is appended to the URL. If `enable-sso=1`, IAM federation attempts to authenticate.
7. Choose **Save** to keep your settings.

To enable service provider-initiated IAM federation IdP

1. Make sure your IAM federation settings are configured and tested. If you're not sure about the configuration, test the connection by using the URLs from the previous procedure.
2. Open Amazon Quick, and choose **Manage Amazon Quick** from your profile menu.
3. Choose **Single sign-on (IAM federation)** from the navigation pane.

4. For **Status**, choose **ON**.
5. Verify that it's working by disconnecting from your IdP and opening Amazon Quick.

To disable service provider initiated IAM federation

1. Open Amazon Quick, and choose **Manage Amazon Quick** from your profile menu.
2. Choose **Single sign-on (IAM federation)** from the navigation pane.
3. For **Status**, choose **OFF**.

Configuring email syncing for federated users in Quick

Applies to: Enterprise Edition

Intended audience: System administrators and Amazon Quick administrators

Note

IAM identity federation doesn't support syncing identity provider groups with Amazon Quick.

In Amazon Quick Enterprise edition, as an administrator you can restrict new users from using personal email addresses when provisioning through their identity provider (IdP) directly to Quick. Quick then uses the preconfigured email addresses passed through the IdP when provisioning new users to your account. For example, you can make it so that only corporate-assigned email addresses are used when users are provisioned to your Amazon Quick account through your IdP.

Note

Make sure that your users are federating directly to Amazon Quick through their IdP. Federating to the AWS Management Console through their IdP and then clicking into Amazon Quick results in an error and they won't be able to access Amazon Quick.

When you configure email syncing for federated users in Amazon Quick, users who log in to your Amazon Quick account for the first time have preassigned email addresses. These are used to register their accounts. With this approach, users can manually bypass by entering an email address. Also, users can't use an email address that might differ from the email address prescribed by you, the administrator.

Amazon Quick supports provisioning through an IdP that supports SAML or OpenID Connect (OIDC) authentication. To configure email addresses for new users when provisioning through an IdP, you update the trust relationship for the IAM role that they use with `AssumeRoleWithSAML` or `AssumeRoleWithWebIdentity`. Then you add a SAML attribute or OIDC token in their IdP. Last, you turn on email syncing for federated users in Amazon Quick.

The following procedures describe these steps in detail.

Step 1: Update the trust relationship for the IAM role with `AssumeRoleWithSAML` or `AssumeRoleWithWebIdentity`

You can configure email addresses for your users to use when provisioning through your IdP to Amazon Quick. To do this, add the `sts:TagSession` action to the trust relationship for the IAM role that you use with `AssumeRoleWithSAML` or `AssumeRoleWithWebIdentity`. By doing this, you can pass in `principal` tags when users assume the role.

The following example illustrates an updated IAM role where the IdP is Okta. To use this example, update the Federated Amazon Resource Name (ARN) with the ARN for your service provider. You can replace items in red with your AWS and IdP service-specific information.

```
{
  "Version": "2012-10-17"      ,
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Federated": "arn:aws:iam::account-id:saml-provider/Okta"
      },
      "Action": "sts:AssumeRoleWithSAML",
      "Condition": {
        "StringEquals": {
          "SAML:aud": "https://signin.aws.amazon.com/saml"
        }
      }
    }
  ],
}
```

```
{
  "Effect": "Allow",
  "Principal": {
    "Federated": "arn:aws:iam::account-id:saml-provider/Okta"
  },
  "Action": "sts:TagSession",
  "Condition": {
    "StringLike": {
      "aws:RequestTag/Email": "*"
    }
  }
}
```

Step 2: Add a SAML attribute or OIDC token for the IAM principal tag in your IdP

After you update the trust relationship for the IAM role as described in the preceding section, add a SAML attribute or OIDC token for the IAM Principal tag in your IdP.

The following examples illustrate a SAML attribute and an OIDC token. To use these examples, replace the email address with a variable in your IdP that points to a user's email address. You can replace items highlighted in red with your information.

- **SAML attribute:** The following example illustrates a SAML attribute.

```
<Attribute Name="https://aws.amazon.com/SAML/Attributes/PrincipalTag:Email"><AttributeValue>john.doe@example.com</AttributeValue></Attribute>
```

Note

If you're using Okta as your IdP, make sure to turn on a feature flag in your Okta user account to use SAML. For more information, see [Okta and AWS Partner to Simplify Access Via Session Tags](#) on the Okta blog.

- **OIDC token:** The following example illustrates an OIDC token example.

```
"https://aws.amazon.com/tags": {"principal_tags": {"Email": ["john.doe@example.com"]}}
```

Step 3: Turn on email syncing for federated users in Amazon Quick

As described preceding, update the trust relationship for the IAM role and add a SAML attribute or OIDC token for the IAM Principal tag in your IdP. Then turn on email syncing for federated users in Amazon Quick as described in the following procedure.

To turn on email syncing for federated users

1. From any page in Amazon Quick, choose your username at top right, and then choose **Manage Amazon Quick**.
2. Choose **Single sign-on (IAM federation)** in the menu at left.
3. On the **Service Provider Initiated IAM federation** page, for **Email Syncing for Federated Users**, choose **ON**.

When email syncing for federated users is on, Amazon Quick uses the email addresses that you configured in steps 1 and 2 when provisioning new users to your account. Users can't enter their own email addresses.

When email syncing for federated users is off, Amazon Quick asks users to input their email address manually when provisioning new users to your account. They can use any email addresses that they want.

Tutorial: Amazon Quick and IAM identity federation

Applies to: Enterprise Edition and Standard Edition

Intended audience: Amazon Quick Administrators and Amazon Quick developers

Note

IAM identity federation doesn't support syncing identity provider groups with Amazon Quick.

In the following tutorial, you can find a walkthrough for setting up the IdP Okta as a federation service for Amazon Quick. Although this tutorial shows the integration of AWS Identity and Access

Management (IAM) and Okta, you can also replicate this solution using your choice of SAML 2.0 IdPs.

In the following procedure, you create an app in the Okta IdP using their "AWS Account Federation" shortcut. Okta describes this integration app as follows:

"By federating Okta to Amazon Web Services (AWS) Identity and Access Management (IAM) accounts, end users get single sign-on access to all their assigned AWS roles with their Okta credentials. In each AWS account, administrators set up federation and configure AWS roles to trust Okta. When users sign in to AWS, they get Okta single sign-in experience to see their assigned AWS roles. They can then select a desired role, which defined their permissions for the duration of their authenticated session. Customers with large numbers of AWS Accounts, check out the AWS Single Sign-On app as an alternative." (<https://www.okta.com/aws/>)

To create an Okta app using Okta's "AWS Account Federation" application shortcut

1. Sign in to your Okta dashboard. If you don't have one, create a free Okta Developer Edition account by using [this Amazon Quick-branded URL](#). When you have activated your email, sign in to Okta.
2. On the Okta website, choose **<> Developer Console** at upper left, and then choose **Classic UI**.
3. Choose **Add Applications**, and choose **Add app**.
4. Enter **aws** for **Search**, and choose **AWS Account Federation** from the search results.
5. Choose **Add** to create an instance of this application.
6. For **Application label**, enter **AWS Account Federation - Amazon Quick**.
7. Choose **Next**.
8. For **SAML 2.0, Default Relay State**, enter **<https://quicksight.aws.amazon.com>**.
9. Open the context (right-click) menu for **Identity Provider metadata**, and choose to save the file. Name the file `metadata.xml`. You need this file in the next procedure.

The contents of the file look similar to the following.

```
<md:EntityDescriptor xmlns:md="urn:oasis:names:tc:SAML:2.0:metadata"
  entityID="http://www.okta.com/exkffz2hATwiVft645d5">
  <md:IDPSSODescriptor WantAuthnRequestsSigned="false"
    protocolSupportEnumeration="urn:oasis:names:tc:SAML:2.0:protocol">
    <md:KeyDescriptor use="signing">
      <ds:KeyInfo xmlns:ds="http://www.w3.org/2000/09/xmldsig#">
        <ds:X509Data>
```

```

        <ds:X509Certificate>

MIIDpjCCAo6gAwIBAgIGAXVjA82hMA0GCSqGSIb3DQEBCwUAMIGTMQswCQYDVQQGEwJVUzETMBEG
.
.      (certificate content omitted)
.
QE/6cRdPQ6v/eaFpUL6Asd6q3sBeq+giRG4=
        </ds:X509Certificate>
    </ds:X509Data>
    </ds:KeyInfo>
</md:KeyDescriptor>
    <md:NameIDFormat>urn:oasis:names:tc:SAML:1.1:nameid-format:emailAddress</
md:NameIDFormat>
    <md:NameIDFormat>urn:oasis:names:tc:SAML:2.0:nameid-format:unspecified</
md:NameIDFormat>
    <md:SingleSignOnService Binding="urn:oasis:names:tc:SAML:2.0:bindings:HTTP-
POST" Location="https://dev-1054988.okta.com/app/amazon_aws/exkffz2hATwiVft645d5/
sso/saml"/>
    <md:SingleSignOnService Binding="urn:oasis:names:tc:SAML:2.0:bindings:HTTP-
Redirect" Location="https://dev-1054988.okta.com/app/amazon_aws/
exkffz2hATwiVft645d5/sso/saml"/>
    </md:IDPSSODescriptor>
</md:EntityDescriptor>

```

10. After you have the XML file saved, scroll to the bottom of the Okta page, and choose **Done**.
11. Keep this browser window open, if possible. You need it later in the tutorial.

Next, you create an identity provider in your AWS account.

To create a SAML provider in AWS Identity and Access Management (IAM)

1. Sign in to the AWS Management Console and open the IAM console at <https://console.aws.amazon.com/iam/>.
2. In the navigation pane, choose **Identity providers**, **Create Provider**.
3. Enter the following settings:
 - **Provider Type** – Choose **SAML** from the list.
 - **Provider Name** – Enter **Okta**.
 - **Metadata Document** – Upload the XML file `manifest.xml` from the previous procedure.
4. Choose **Next Step, Create**.

5. Locate the IdP that you created and choose it to view the settings. Note the **Provider ARN**. You need this to finish the tutorial.
6. Verify that the identity provider is created with your settings. In IAM, choose **Identity providers, Okta** (the IdP you added), **Download metadata**. The file should be the one that you recently uploaded.

Next, you create an IAM role to enable the SAML 2.0 federation to act as a trusted entity in your AWS account. For this step, you need to choose how you want to provision users in Amazon Quick. You can do one of the following:

- Grant permission to the IAM role so that first-time visitors become Amazon Quick users automatically.

To create an IAM role for a SAML 2.0 federation as a trusted entity

1. Sign in to the AWS Management Console and open the IAM console at <https://console.aws.amazon.com/iam/>.
2. In the navigation pane, choose **Roles, Create Role**.
3. For **Select type of trusted entity**, choose the card labeled **SAML 2.0 federation**.
4. For **SAML provider**, select the IdP that you created in the previous procedure, for example Okta.
5. Enable the option **Allow programmatic and AWS Management Console access**.
6. Choose **Next: Permissions**.
7. Paste the following policy into the editor.

In the policy editor, update the JSON with your provider's Amazon Resource Name (ARN).

```
{
  "Version": "2012-10-17"
,
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "sts:AssumeRoleWithSAML",
      "Resource": "arn:aws:iam::111111111111:saml-provider/Okta",
      "Condition": {
        "StringEquals": {
          "saml:aud": "https://signin.aws.amazon.com/saml"
        }
      }
    }
  ]
}
```

```

    }
  }
}
]
}

```

8. Choose **Review policy**.
9. For **Name**, enter **QuicksightOktaFederatedPolicy**, and then choose **Create policy**.
10. Choose **Create policy, JSON** a second time.
11. Paste the following policy into the editor.

In the policy editor, update the JSON with your AWS account ID. It should be the same account ID that you used in the previous policy in the provider ARN.

```

{
  "Version": "2012-10-17"      ,
  "Statement": [
    {
      "Action": [
        "quicksight:CreateReader"
      ],
      "Effect": "Allow",
      "Resource": [
        "arn:aws:quicksight::111111111111:user/${aws:userid}"
      ]
    }
  ]
}

```

You can omit the AWS Region name in the ARN, as shown following.

```
arn:aws:quicksight::111111111111:user/${aws:userid}
```

12. Choose **Review policy**.
13. For **Name**, enter **QuicksightCreateReader**, and then choose **Create policy**.
14. Refresh the list of policies by choosing the refresh icon at right.
15. For **Search**, enter **QuicksightOktaFederatedPolicy**. Choose the policy to enable it ().

If you don't want to use automatic provisioning, you can skip the following step.

To add a Amazon Quick user, use [register-user](#). To add a Amazon Quick group, use [create-group](#). To add users to the Amazon Quick group, use [create-group-membership](#).

16. (Optional) For **Search**, enter **QuicksightCreateReader**. Choose the policy to enable it ().

Do this step if you want to provision Amazon Quick users automatically, rather than using the Amazon Quick API.

The QuicksightCreateReader policy activates automatic provisioning by allowing use of the quicksight:CreateReader action. Doing this grants dashboard subscriber (reader-level) access to first-time users. A Amazon Quick administrator can later upgrade them from the Amazon Quick profile menu, **Manage Amazon Quick**, **Manage users**.

17. To continue attaching the IAM policy or policies, choose **Next: Tags**.
18. Choose **Next: Review**.
19. For **Role name**, enter **QuicksightOktaFederatedRole**, and choose **Create role**.
20. Verify that you completed this successfully by taking these steps:
 - a. Return to the main page of the IAM console at <https://console.aws.amazon.com/iam/>. You can use your browser's **Back** button.
 - b. Choose **Roles**.
 - c. For **Search**, enter Okta. Choose **QuicksightOktaFederatedRole** from the search results.
 - d. On the **Summary** page for the policy, examine the **Permissions** tab. Verify that the role has the policy or policies that you attached to it. It should have QuicksightOktaFederatedPolicy. If you chose to add the ability to create users, it should also have QuicksightCreateReader.
 - e. Use the  icon to open each policy. Verify that the text matches what is shown in this procedure. Double-check that you added your own AWS account number in place of the example account number 111111111111.
 - f. On the **Trust relationships** tab, verify that the **Trusted entities** field contains the ARN for the identity provider. You can double-check the ARN in the IAM console by opening **Identity providers**, **Okta**.
 - g.

To create an access key for Okta

1. Sign in to the AWS Management Console and open the IAM console at <https://console.aws.amazon.com/iam/>.
2. Add a policy that allows Okta to display a list of IAM roles to the user. To do this, choose **Policy, Create policy**.
3. Choose **JSON**, then enter the following policy.

```
{
  "Version": "2012-10-17"      ,
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "iam:ListRoles",
        "iam:ListAccountAliases"
      ],
      "Resource": "*"
    }
  ]
}
```

4. Choose **Review Policy**.
5. For **Name**, enter **OktaListRolesPolicy**. Then choose **Create policy**.
6. Add a user so you can provide Okta with an access key.

In the navigation pane, choose **Users, Add User**.

7. Use the following settings:
 - For **User name**, enter OktaSSOUser.
 - For **Access type**, enable **Programmatic access**.
8. Choose **Next: Permissions**.
9. Choose **Attach existing policies directly**.
10. For **Search**, enter **OktaListRolesPolicy**, and choose **OktaListRolesPolicy** from the search results.
11. Choose **Next: Tags**, and then choose **Next: Review**.
12. Choose **Create user**. Now you can get the access key.

13. Download the key file by choosing **Download .csv**. The file contains the same access key ID and secret access key that displays on this screen. However, because AWS doesn't display this information a second time, make sure to download the file.
14. Verify that you completed this step correctly by doing the following:
 - a. Open the IAM console, and choose **Users**. Search for **OktaSSOUser**, and open it by choosing the username from the search results.
 - b. On the **Permissions** tab, verify that the **OktaListRolesPolicy** is attached.
 - c. Use the  icon to open the policy. Verify that the text matches what is shown in this procedure.
 - d. On the **Security credentials** tab, you can check the access key, although you already downloaded it. You can return to this tab to create an access key when you need a new one.

In the following procedure, you return to Okta to provide the access key. The access key works with your new security settings to allow AWS and the Okta IdP to work together.

To finish configuring the Okta application with AWS settings

1. Return to your Okta dashboard. If requested to do so, sign in. If the developer console is no longer open, choose **Admin** to reopen it.
2. If you have to reopen Okta, you can return to this section by following these steps:
 - a. Sign in to Okta. Choose **Applications**.
 - b. Choose **AWS Account Federation - Amazon Quick**—the application that you created at the beginning of this tutorial.
 - c. Choose the **Sign On** tab, between **General** and **Mobile**.
3. Scroll to **Advanced Sign-On Settings**.
4. For **Identity Provider ARN (Required only for SAML IAM federation)**, enter the provider ARN from the previous procedure, for example:

`arn:aws:iam::111122223333:saml-provider/Okta`
5. Choose **Done** or **Save**. The name of the button varies depending if you are creating or editing the application.

6. Choose the **Provisioning** tab, and at the lower part of the tab, choose **Configure API Integration**.
7. Turn on **Enable API integration** to display the settings.
8. For **Access Key** and **Secret Key**, provide the access key and secret key that you downloaded previously to a file named **OktaSSOUser_credentials.csv**.
9. Choose **Test API Credentials**. Look above the **Enable API integration** setting for a message confirming that **AWS Account Federation was verified successfully**.
10. Choose **Save**.
11. Make sure that **To App** is highlighted at left, and choose **Edit** at right.
12. For **Create Users**, turn on the option **Enable**.
13. Choose **Save**.
14. On the **Assignments** tab, near **Provisioning** and **Import**, choose **Assign**.
15. Do one or more of the following to enable federated access:
 - To work with individual users, choose **Assign to People**.
 - To work with IAM groups, choose **Assign to Groups**. You can choose specific IAM groups or **Everyone (All users in your organization)**.
16. For each IAM user or group, do the following:
 - a. Choose **Assign, Role**.
 - b. Select **QuicksightOktaFederatedRole** from the list of IAM roles.
 - c. For **SAML User Roles**, enable **QuicksightOktaFederatedRole**.
17. Choose **Save and Go Back**, and then choose **Done**.
18. Verify that you completed this step correctly by choosing the **People** or **Groups** filter at left, and checking the users or groups that you entered. If you can't complete this process because the role that you created doesn't appear in the list, return to the previous procedures to verify the settings.

To sign in to Amazon Quick using Okta (IdP to service provider sign-in)

1. If you are using an Okta administrator account, switch to user mode.
2. Sign in to your Okta Applications dashboard with a user that has been granted federated access. You should see a new application with your label, for example **AWS Account Federation - Amazon Quick**.

3. Choose the application icon to launch **AWS Account Federation - Amazon Quick**.

You can now manage identities using Okta and use federated access with Quick.

The following procedure is an optional part of this tutorial. If you follow its steps, you authorize Amazon Quick to forward authorizations requests to the IdP on behalf of your users. Using this method, users can sign in to Amazon Quick with no need to sign in using the IdP page first.

(Optional) To set up Amazon Quick to send authentication requests to Okta

1. Open Amazon Quick, and choose **Manage Amazon Quick** from your profile menu.
2. Choose **Single sign-on (IAM federation)** from the navigation pane.
3. For **Configuration**, **IdP URL**, enter the URL that your IdP provides to authenticate users, for example `https://dev-1-----0.okta.com/home/amazon_aws/0oabababababaGQei5d5/282`. You can find this in your Okta app page, on the **General** tab, in **Embed Link**.
4. For **IdP URL**, enter `RelayState`.
5. Do one of the following:
 - To test signing in with your identity provider first, use the custom URL provided in **Test starting with your IdP**. You should arrive at the start page for Amazon Quick, for example `https://quicksight.aws.amazon.com/sn/start`.
 - To test signing in with Amazon Quick first, use the custom URL provided in **Test the end-to-end experience**. The `enable-sso` parameter is appended to the URL. If `enable-sso=1`, IAM federation attempts to authenticate. If `enable-sso=0`, Amazon Quick doesn't send the authentication request, and you sign in to Amazon Quick as before.
6. For **Status**, choose **ON**.
7. Choose **Save** to keep your settings.

You can create a deep link to a Amazon Quick dashboard to allow users to use IAM federation to connect directly to specific dashboards. To do this, you append the relay state flag and dashboard URL to the Okta single sign-on URL, as described following.

To create a deep link to a Amazon Quick dashboard for single sign-on

1. Locate the Okta application's single sign-on (IAM federation) URL in the `metadata.xml` file that you downloaded beginning of the tutorial. You can find the URL near the bottom of the

file, in the element named `md:SingleSignOnService`. The attribute is named `Location` and the value ends with `/sso/saml`, as shown in the following example.

```
<md:SingleSignOnService Binding="urn:oasis:names:tc:SAML:2.0:bindings:HTTP-Redirect" Location="https://dev-0000001.okta.com/app/amazon_aws/abcdef2hATwiVft645d5/sso/saml"/>
```

2. Take the value of the IAM federation URL and append `?RelayState=` followed by the URL of your Amazon Quick dashboard. The `RelayState` parameter relays the state (the URL) that the user was in when they were redirected to the authentication URL.
3. To the new IAM federation with the relay state added, append the URL of your Amazon Quick dashboard. The resulting URL should resemble the following.

```
https://dev-1-----0.okta.com/app/amazon_aws/abcdef2hATwiVft645d5/sso/saml?RelayState=https://us-west-2.quicksight.aws.amazon.com/sn/analyses/12a12a2a-121a-212a-121a-abcd12abc1ab
```

4. If the link you create doesn't open, check that you are using the most recent IAM federation URL from the `metadata.xml`. Also check that the username you use to sign in isn't assigned in more than one IAM federation Okta app.

Using Active Directory with Amazon Quick Enterprise edition

Applies to: Enterprise Edition

Intended audience: System administrators

Note

IAM identity federation doesn't support syncing identity provider groups with Amazon Quick.

Amazon Quick Enterprise edition supports both [AWS Directory Service for Microsoft Active Directory](#) and [Active Directory Connector](#).

To create a new directory to be your identity manager for Quick, use AWS Directory Service for Microsoft Active Directory, also known as AWS Managed Microsoft AD. This is an Active Directory host in the AWS Cloud that offers most of the same functionality of Active Directory. Currently, you can connect to Active Directory in any AWS Region supported by Amazon Quick, except for Asia Pacific (Singapore). When you create a directory, you use it with a virtual private cloud (VPC). For more information, see [VPC](#).

If you have an existing directory that you want to use for Quick, you can use Active Directory Connector. This service redirects directory requests to your Active Directory—in another AWS Region or on-premises—without caching any information in the cloud.

For a walkthrough about creating and managing a directory with AWS Managed Microsoft AD, see [Use an AWS Managed Microsoft AD with Quick?](#) in the AWS Knowledge Center.

When you use AWS Directory Service to launch a directory, AWS creates an organizational unit (OU) with the same name as your domain. AWS also creates an administrative account with delegated administrative rights for the OU. You can create accounts, groups, and policies within the OU by using Active Directory users and groups. For more information, see [Best Practices for AWS Managed Microsoft AD](#) in the *Directory Service Administration Guide*.

After you establish your directory, you use it with Quick by creating groups for users. Amazon Quick has six specific user roles that can be assigned, including Pro versions that provide access to advanced capabilities:

- **Quick admins** – Admins can change account settings, manage accounts. Admins can also purchase additional Amazon Quick user subscriptions or [SPICE](#) capacity, or cancel the subscription to Amazon Quick for your AWS account. Admin Pro users have additional capabilities including creating content using natural language, building knowledge bases, configuring actions, and accessing advanced automation workflows.
- **Quick authors** – Amazon Quick authors can create data sources, datasets, analyses, and dashboards. They can share analyses and dashboards with other Amazon Quick users. Author Pro users can additionally create content using natural language, build knowledge bases, configure actions, and access advanced automation capabilities.
- **Quick readers** – Readers can view and interact with dashboards that were created by someone else. Reader Pro users have access to advanced features including AI chat agents, collaborative spaces, flows, and extensions.

You can add or refine access by applying IAM policies. For example, you can use IAM policies to allow users to subscribe themselves.

When you subscribe to Amazon Quick Enterprise edition and choose Active Directory as your identity provider, you can associate your AD groups with Amazon Quick. You can also add or change your AD groups later on.

Topics

- [Directory integration with Quick Enterprise edition](#)

Directory integration with Quick Enterprise edition

Applies to: Enterprise Edition

Intended audience: System administrators

Note

IAM identity federation doesn't support syncing identity provider groups with Amazon Quick.

Quick Enterprise supports the following options:

- AWS Directory Service
- AWS Directory Service with AD Connector
- On-premises Active Directory with IAM federation or AD Connector
- IAM federation using AWS IAM Identity Center or another third-party federation service

If you want to use IAM federation with an on-premises Active Directory, you implement AWS Directory Service as a separate Active Directory with a trust relationship to the on-premises Active Directory.

If you want to avoid using a trust relationship, you can deploy a standalone domain for authentication within AWS. Then you can create users and groups in Active Directory. You'd then

map them to users and groups in Quick. In this example, users authenticate using their Active Directory login credentials. To make access to Quick transparent to your users, use IAM federation in this scenario.

Using multi-factor authentication (MFA) with Amazon Quick

Applies to: Enterprise Edition and Standard Edition

Intended audience: System administrators

Important

Amazon Quick recommends that you integrate new Quick subscriptions with IAM Identity Center for identity management. This IAM identity federation user guide is provided as a reference for existing account configurations. For more information on integrating your Quick account with IAM Identity Center, see [Configure your Quick account with IAM Identity Center](#).

Note

IAM identity federation doesn't support syncing identity provider groups with Amazon Quick.

There are several ways that you can use multi-factor authentication (MFA) with Quick. You can use it with AWS Identity and Access Management (IAM). You can use it with AD Connector or your [AWS Directory Service](#) for Microsoft Active Directory, also known as AWS Microsoft Active Directory or AWS Managed Microsoft Active Directory. And if you use an external identity provider (IdP), AWS doesn't need to have any information about MFA because that is part of the authentication handled by the IdP.

For more information, see the following:

- [Using multi-factor authentication \(MFA\) in AWS](#) in the IAM User Guide

- [Enable Multi-Factor Authentication for AWS Managed Microsoft AD](#) in the AWS Directory Service Administration Guide
- [Enable Multi-Factor Authentication for AD Connector](#) in the AWS Directory Service Administration Guide

If you're a developer, see the following:

- [How do I use an MFA token to authenticate access to my AWS resources through the AWS CLI](#) in the [AWS Knowledge Center](#)
- [Configuring MFA-protected API access](#) in the IAM User Guide

Allowlisting Amazon Quick domains

If your end users are signing in to Amazon Quick using AWS root (not recommended), AWS Identity and Access Management (IAM), corporate Active Directory, or native Quick credentials, make sure to allow-list the following domains within your organization's network.

User type	Domain or domains to allow-list
Users who sign in directly through Amazon Quick and Active Directory users	signin.aws and awsapps.com
AWS root user	signin.aws.amazon.com and amazon.com
IAM users	signin.aws.amazon.com

Important

We strongly recommend that you don't use the AWS root user for your everyday tasks, even the administrative ones. Instead, adhere to the best practice of using the root user only to create your first IAM user. Then securely lock away the root user credentials and use them to perform only a few account and service management tasks. For more information, see [AWS account root user](#) in the *IAM User Guide*.

Configuring VPC connections in Amazon Quick Sight

Applies to: Enterprise Edition

Intended audience: System administrators and Amazon Quick administrators

Note

If you're a Amazon Quick system administrator configuring a VPC connection to Amazon Quick Sight, this section is for you. Amazon Quick knowledge bases currently don't support VPC integrations.

Quick Enterprise edition is fully integrated with the Amazon VPC service. A VPC based on this service closely resembles a traditional network that you operate in your own data center. It enables you to secure and isolate traffic between resources. You define and control the network elements to suit your requirements, while still getting the benefit of cloud networking and the scalable infrastructure of AWS.

By creating a VPC connection in Amazon Quick, you're adding elastic network interfaces in your VPC. These network interfaces allow Amazon Quick to exchange network traffic with a network instance within your VPC. You can provide all of the standard security controls for this network traffic, as you do with other traffic in your VPC. Route tables, network access control lists (ACLs), subnets, and security groups settings all apply to network traffic to and from Amazon Quick in the same way that they apply to traffic between other instances in your VPC.

When you register a VPC connection with Amazon Quick, you can securely connect to data that's available only in your VPC, for example:

- Data you can reach by IP address
- Data that isn't available on the public internet
- Private databases
- On-premises data

This works if you set up connectivity between the VPC and your on-premises network. For example, you might set up connectivity with AWS Direct Connect, a virtual private network (VPN), or a proxy.

After you connect to the data, you can use it to create data analyses and publish secure data dashboards.

To further increase security, consider logging data access operations with AWS CloudTrail, as described in [Logging Amazon Quick information with CloudTrail](#). You can even create a dashboard to help you analyze your CloudTrail logs. By combining Amazon Quick logs with logs from your other AWS services, you can get a fuller view of how your data is being used.

You don't need to be a networking expert to connect and use a VPC with Amazon Quick, because Amazon Quick provides a user interface for adding your network information. However, the person who gathers the information that you need for setup should have some understanding of networking concepts and using VPCs. This person also needs read-only access to the services. If network changes are required, we recommend that you don't make changes to your networking configuration without expert assistance.

To use a command line interface to access your VPC, you can use the AWS Command Line Interface (AWS CLI). For more information on using the AWS CLI, see the [AWS CLI User Guide](#).

Topics

- [VPC terminology](#)
- [Supported VPC data sources](#)
- [Setting up a VPC to use with Amazon Quick](#)
- [Finding information to connect to a VPC](#)

VPC terminology

The following terminology can be useful when you work with a VPC and Amazon Quick.

A *VPC* is a virtual private cloud, which works like a private network to isolate the resources within it. The solution described in these topics uses an AWS service called Amazon VPC.

A *route table* contains a set of rules, called *routes*, that are used to determine where network traffic is directed. You can view the route table in the Amazon VPC console at <https://>

console.aws.amazon.com/vpc/. The VPC details display the route table that the VPC is using. You can also see **Route tables** listed in the Amazon VPC console.

A *subnet* is a defined set of network IP addresses that are used to increase the security and efficiency of network communications. You can think of them like postal codes, used for routing packages from one location to another. The **Subnets** list in the Amazon VPC console displays subnet IDs and also their associated VPC IDs, route tables, and network ACLs. You need to provide at least two subnets in different availability zones to create a VPC connection.

A *network interface* represents a virtual network card. The network interface automatically created by Amazon Quick is called a *Amazon Quick network interface*. Each network interface in a VPC connection is configured based on the subnet it's attached to. You can view your Amazon Quick network interfaces in the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>. The network interface displays its network interface ID, subnet ID, VPC ID, security group, and the Availability Zone that it exists in. You can click on the security group name to see its group ID and its inbound and outbound rules. The term *network interface* in the following sections always means elastic network interface.

A *security group* is a set of rules that controls the network access to the resources it is associated with. Access is permitted only to and from the components defined in the security group's inbound and outbound rules. If no rules are defined, the security group prevents all access. You can view security groups from several different consoles, depending on which resource that a particular security group applies to. You can see all the security groups and their settings in one place in the VPC console. For the Amazon Quick VPC connection, create a new security group.

Inbound and outbound rules define the following:

- The type of traffic to allow, for example "**All TCP**" or "**RDS**".
- The protocol to allow (TCP, UDP, or ICMP).
- The traffic source to allow for inbound rules, or the traffic destination to allow for outbound rules. When you work with a VPC and Amazon Quick, you specify the security group ID to use.
- An optional description. We recommend that you add the word **Amazon Quick** to the description for Amazon Quick VPC rules.

An *internet gateway* is a VPC component that allows communication between instances in your VPC and the internet. You don't need an internet gateway to use Amazon Quick VPC connections.

A *VPC endpoint* enables you to privately connect your VPC to supported AWS services without using public IP addresses. You don't need to set up a VPC endpoint to use Amazon Quick VPC connections.

Supported VPC data sources

With a VPC connection, you can connect Amazon Quick to a data source inside your VPC over a private network path. Traffic stays off the public internet. Not every data source can use a VPC connection. VPC connections support the following two categories:

Amazon Quick Sight data sources

Structured data sources, such as databases and data warehouses, that you connect to when you create datasets and analyses. Most of these support VPC connections. For more information about supported Amazon Quick Sight data sources, see [Amazon Quick Sight data sources that support VPC connections](#).

Knowledge base data sources

Data sources that you connect to when you create a knowledge base, including document repositories and content management systems. Amazon Redshift supports VPC connections in this category. For more information about knowledge base data sources, see [Knowledge base data sources that support VPC connections](#).

Both categories share the same configuration requirements. For more information about those requirements, see [Requirements for a VPC data source](#).

Amazon Quick Sight data sources that support VPC connections

The following Amazon Quick Sight data sources support VPC connections. When you create a dataset from one of these sources, you can route the connection through a VPC connection:

- Amazon OpenSearch Service
- Amazon Redshift
- Amazon Relational Database Service
- Amazon Aurora
- Apache Impala
- Databricks

- Exasol
- MariaDB
- Microsoft SQL Server
- MySQL
- Oracle
- PostgreSQL
- Presto
- Snowflake
- Starburst Enterprise
- Teradata
- Trino

Other Amazon Quick Sight data sources can't use a VPC connection. This group includes Amazon Athena, Google BigQuery, Apache Spark, and Amazon Timestream. It also includes the connectors for software as a service (SaaS) applications, such as Salesforce, ServiceNow, Jira, and GitHub. If you attach a VPC connection to a data source that doesn't support one, the request fails with an invalid parameter error.

Knowledge base data sources that support VPC connections

Amazon Redshift supports VPC connections when you use it as a knowledge base data source. When you create a knowledge base from an Amazon Redshift data source, you can choose a VPC connection. Your queries then reach the cluster over a private network path.

The following knowledge base data sources can't use a Quick VPC connection:

- Amazon Bedrock managed knowledge bases
- Amazon Q Business indexes
- Amazon S3
- Box
- Confluence
- Google Drive
- Microsoft OneDrive

- Microsoft SharePoint

If you attach a VPC connection to one of these data sources, the request fails with an invalid parameter error.

 Note

You can still restrict Amazon S3 bucket access to a VPC endpoint. That is a separate mechanism, and you control it through your bucket policy. For more information about restricting Amazon S3 bucket access to a VPC endpoint, see [Configure VPC access for the Amazon S3 connector](#).

Requirements for a VPC data source

Before you can reach a data source in your VPC from Amazon Quick, your configuration must meet the following requirements:

1. The Domain Name System (DNS) name of the VPC data source must resolve from outside your VPC.
2. The connection must return the private IP address of your instance. Databases hosted by Amazon Redshift, Amazon RDS, and Aurora automatically meet this requirement.
3. A clearly defined network path must connect the data source and Amazon Quick.
4. A VPC connection for the VPC must exist in Amazon Quick. You create or select one in the Amazon Quick console.

Setting up a VPC to use with Amazon Quick

Applies to: Enterprise Edition

Intended audience: System administrators

To set up a VPC to use with Amazon Quick Enterprise edition, you need access to Amazon VPC and Amazon EC2. You also need access to each AWS database service that you plan to add to Quick.

You can use the console, or you can use the AWS Command Line Interface (AWS CLI). For more information about the CLI, see the [AWS Command Line Interface User Guide](#). To work with the CLI, go to <https://aws.amazon.com/cli/>.

Before you begin to set up your VPC connection in Amazon Quick, make sure that you understand the components of a VPC deployment. As part of that, familiarize yourself with the VPC's subnets and security groups in relation to the destinations (databases) that you want to reach from Amazon Quick. To set up a successful VPC connection, make sure that the following components work together to allow network traffic to pass between Amazon Quick and your data source:

- The Amazon VPC service
- The subnets that your data source is using
- The Amazon Quick elastic network interfaces and the subnets they use
- The route table
- Inbound and outbound rules for these security groups:
 - Security group for your VPC. We recommend you create a new security group to isolate the rules on the VPC security group from the rules on the Amazon Quick network interface's security group).
 - Security group attached to the Amazon Quick network interface.
 - Security group attached to the database server (for each database server that you want to use).
- (Optional) Amazon Route 53 Resolver inbound endpoints for private DNS resolution.

In the following topics, you can find the network components that are involved. You can also find descriptions of their roles in the network configuration of your VPC and your Amazon Quick VPC connection. The network interface for Amazon Quick that is automatically created during setup is called the *Amazon Quick network interface (QNI)*.

If your VPC is already completely configured, skip to the next section, [Finding information to connect to a VPC](#).

Topics

- [VPC](#)
- [Subnets](#)
- [Security groups: inbound and outbound rules](#)

- [Sample rules](#)
- [Route table](#)
- [Amazon Quick elastic network interface](#)
- [Inbound endpoints for Amazon Route 53 Resolver](#)

VPC

A *virtual private cloud (VPC)* is a virtual network dedicated to your AWS account. The Amazon VPC service that provides it is a networking layer for your AWS resources. Using Amazon VPC, you can define a virtual network in your own logically isolated area within the AWS Cloud. A VPC closely resembles a traditional network that you might operate in your own data center, with the benefits of using the AWS scalable infrastructure. Amazon VPC for Amazon EC2 virtual computing environments, known as *instances*, can be used for a variety of AWS resources.

VPCs offer options that allow for flexibility in a secure environment, for example:

- To configure your VPC, you can set its IP address range, create subnets, configure route tables, network gateways, network interfaces, and security settings.
- To make the AWS Cloud an extension of your data center, you can connect your VPC to your own corporate data center.
- You can connect your instances in the VPC to the internet, or keep your instances isolated on a private network.
- To protect the resources in each subnet, you can use multiple layers of security, including security groups and network access control lists (ACLs).

For more information, see the [Amazon VPC User Guide](#).

If you have a default VPC and don't specify a subnet when you launch an instance, the instance is launched into your default VPC. You can launch instances into your default VPC without needing to know anything about Amazon VPC.

If you don't already have a VPC or want to use a new one, you can create one by following the instructions in [Getting started with Amazon VPC](#) in the *Amazon VPC User Guide*. This section offers guidance on how to set up your VPC. The guidance includes options for public and private subnets and for AWS Site-to-Site VPN access for your corporate network (known as *on-premises access*). You can also use VPC peering or Direct Connect to reach an on-premises database instance.

Using the AWS CLI

You can start to set up a VPC in Amazon EC2 by using the [aws ec2 create-vpc](#) command. To learn more about VPC settings for the AWS CLI, see [Examples for VPC](#) in the *Amazon VPC User Guide*.

Using the Amazon EC2 console

To view your VPC or create a new one in Amazon EC2, sign in to the AWS Management Console and open the Amazon VPC console at <https://console.aws.amazon.com/vpc/>. To create a new VPC, choose **Launch VPC Wizard** and follow the instructions. Note your new VPC ID for future reference. To view VPCs, choose **Your VPCs** on the left side.

Amazon VPC resources in VPC guides and AWS Support articles

For general information, see [Working with VPCs and subnets](#).

For step-by-step instructions for setting up a VPC, see the following topics (choose the ones that relate to your scenario):

- [Create an IPv4 VPC and subnets using the AWS CLI](#)
- [Sharing public subnets and private subnets](#)
- [Working with site-to-site VPN](#)
- [AWS Site-to-Site VPN Network Administrator Guide](#) (choose your network device for specific instructions)
- [Generic Customer Gateway Device Without Border Gateway Protocol](#) (recommended for customer gateways)

If you want to migrate data source instances into the same VPC, see the following AWS Support articles:

- [How do I change the VPC for an Amazon RDS DB instance?](#)
- [How do I move my EC2 instance to another subnet, Availability Zone, or VPC?](#)
- [How do I move my Amazon Redshift cluster from one VPC to another VPC?](#)

For troubleshooting information, see [How do I troubleshoot issues with VPC route tables?](#), an article with video created by AWS Support.

Subnets

A *subnet* is a range of IP addresses in your VPC. You need to provide at least two subnets to create a VPC connection. Each subnet must belong a different availability zone. You can attach AWS resources, such as Amazon EC2 instances and Amazon RDS DB instances, to subnets. You can create subnets to group instances together according to your security and operational needs.

For Amazon Quick to connect to your database, the network needs to route traffic to the data sources that you want to reach from one of the subnets used by the Amazon Quick network interface. Amazon Quick determines which subnet to route traffic through on the backend. If the availability zone that the subnet is attached to experiences an outage, Amazon Quick reroutes the traffic to one of the other subnets that are configured in the VPC connection. If the data sources are on different subnets, make sure that there is a route from the Amazon Quick network interface to your database instance. By default, each subnet in a VPC is associated with one main route table and can reach the other subnets. For more information, see [VPC and Subnets](#) and [Network ACLs](#) in the *Amazon VPC User Guide*.

If you use Amazon RDS, DB instances are associated with a subnet group that you can view either in the Amazon RDS console (<https://console.aws.amazon.com/rds/>) or in the VPC console. For troubleshooting connectivity to Amazon RDS, see the AWS Support article [How can I troubleshoot connectivity to an Amazon RDS instance that uses a public or private subnet of a VPC?](#)

Security groups: inbound and outbound rules

A *security group* acts as a virtual firewall for your instance to control inbound and outbound traffic. For each security group, you add rules that control the inbound traffic to instances, and a separate set of rules that control the outbound traffic.

For your VPC connection, create a new security group with the description QuickSight-VPC. This security group must allow all inbound TCP traffic from the security groups of the data destinations that you want to reach. The following example creates a new security group in the VPC and returns the ID of the new security group.

```
aws ec2 create-security-group \  
--group-name quicksight-vpc \  
--description "QuickSight-VPC" \  
--vpc-id vpc-0daeb67adda59e0cd
```


⚠ Important

Network configuration is sufficiently complex that we strongly recommend that you create a new security group for use with Amazon Quick. It also makes it easier for AWS Support to help you if you need to contact them. Creating a new group isn't absolutely required. However, the following topics are based on the assumption that you follow this recommendation.

To enable Quick to successfully connect to an instance in your VPC, configure your security group rules to allow traffic between the Amazon Quick network interface and the instance that contains your data. To do this, configure the security group attached to your database's instance inbound rules to allow the following traffic:

- From the port that Amazon Quick is connecting to
- From one of the following options:
 - The security group ID that's associated with Amazon Quick network interface (recommended)
 - or
 - The private IP address of the Amazon Quick network interface

For more information, see [Security groups for your VPC](#) and [VPCs and subnets](#) in the *Amazon VPC User Guide*.

Use the topics listed below to learn more about inbound and outbound rules.

Topics

- [Inbound rules](#)
- [Outbound rules](#)

Inbound rules**⚠ Important**

The following section applies to your VPC connection if the connection was created before April 27, 2023.

When you create a security group, it has no inbound rules. No inbound traffic originating from another host to your instance is allowed until you add inbound rules to the security group.

The security group attached to the Amazon Quick network interface behaves differently than most security groups, because it isn't stateful. Other security groups are usually *stateful*. This means that, after they establish an outbound connection to a resource's security group, they automatically allow return traffic. In contrast, the Amazon Quick network interface security group doesn't automatically allow return traffic. Because of this, adding an egress rule to the Amazon Quick network interface security group doesn't work. To make it work for the Amazon Quick network interface security group, make sure to add an inbound rule that explicitly authorizes the return traffic from the database host.

The inbound rule in your security group must allow traffic on all ports. It needs to do this because the destination port number of any inbound return packets is set to a randomly allocated port number.

To restrict Amazon Quick to connect only to certain instances, you can specify the security group ID (recommended) or private IP address of the instances that you want to allow. In either case, your security group inbound rule still needs to allow traffic on all ports (0–65535).

To allow Amazon Quick to connect to any instance in the VPC, you can configure the Amazon Quick network interface security group. In this case, give it an inbound rule to allow traffic on 0.0.0.0/0 on all ports (0–65535). The security group used by the Amazon Quick network interface should be different than the security groups used for your databases. We recommend that you use separate security groups for VPC connection.

Important

If you are using a long-standing Amazon RDS DB instance, check your configuration to see if you're using a DB security group. DB security groups are used with DB instances that are not in a VPC and are on the EC2-Classic platform.

If this is your configuration, and you aren't moving your DB instance into the VPC for use with Amazon Quick, make sure to update your DB security group's inbound rules. Update them to allow inbound traffic from the VPC security group that you're using for Amazon Quick. For more information, see [Controlling Access with Security Groups](#) in the *Amazon RDS User Guide*.

Outbound rules

Important

The following section applies to your VPC connection if the connection was created before April 27, 2023.

By default, a security group includes an outbound rule that allows all outbound traffic. We recommend that you remove this default rule and add outbound rules that allow specific outbound traffic only.

Warning

Do not configure the security group on the Amazon Quick network interface with an outbound rule to allow traffic on all ports. For information on key considerations and recommendations for managing network egress traffic from VPCs, see [Security best practices for your VPC](#) in the *Amazon VPC User Guide*.

The security group attached to Amazon Quick network interface should have outbound rules that allow traffic to each of the database instances in your VPC that you want Amazon Quick to connect to. To restrict Amazon Quick to connect only to certain instances, specify the security group ID (recommended) or the private IP address of the instances to allow. You set this up, along with the appropriate port numbers for your instances (the port that the instances are listening on), in the outbound rule.

The VPC security group must also allow outbound traffic to the security groups of the data destinations, specifically on the port or ports that the database is listening on.

Sample rules

Following, you can find some example configurations of inbound and outbound rules for Amazon RDS and Amazon Redshift.

VPC connection rules: Amazon Quick Sight: Amazon RDS for MySQL

The following tables show rule settings for connecting Amazon Quick Sight to Amazon RDS for MySQL.

Amazon Quick Sight Network interface security group: inbound rule

Type	All TCP
Protocol	TCP
Port Range	0 - 65535
Source	<i>sg-RDS11111111</i>
Description	Amazon Quick Sight - RDS MySQL

Amazon Quick Sight Network interface security group: outbound rule

Type	MYSQL/Aurora
Protocol	TCP
Port Range	3306
Source	sg-RDS11111111
Description	Amazon Quick Sight to RDS MySQL

RDS MySQL: inbound rule

Type	MYSQL/Aurora
Protocol	TCP
Port Range	3306
Source	sg-ENI3333333
Description	Amazon Quick Sight to RDS MySQL

VPC connection rules: Amazon Redshift in Amazon Quick Sight

The following tables show rule settings for connecting Amazon Quick Sight to Amazon Redshift.

Amazon Quick Sight network interface security group: inbound rule

Type	All TCP
Protocol	TCP
Port Range	0 - 65535
Source	sg-RedSh222222
Description	Amazon Quick Sight– Amazon Redshift

Amazon Quick Sight network interface security group: outbound rule

Type	Amazon Redshift
Protocol	TCP
Port Range	5439
Source	sg-RedSh222222
Description	Amazon Quick Sight– Amazon Redshift

Amazon Redshift: inbound rule

Type	Amazon Redshift
Protocol	TCP
Port Range	5439
Source	sg-ENI3333333

Description	Amazon Quick Sight– Amazon Redshift
-------------	--

Route table

To use VPC peering or Direct Connect to reach an on-premises database instance, update the route table that's associated with the VPC you're using with Amazon Quick. For more information on route tables, see [Route tables](#) in the *Amazon VPC User Guide*.

To learn more about VPC peering and view sample scenarios and configurations, see [What is VPC peering?](#) in the *Amazon VPC Peering Guide*. For an example configuration, see [Example: Services using AWS PrivateLink and VPC peering](#) in the *Amazon VPC User Guide*.

Using the AWS CLI

The following example creates a route table.

```
aws ec2 create-route-table --vpc-id vpc-0daeb67adda59e0cd
```

Then you can use the `create-route` command to create a route. For more information and examples, see [create-route](#) in the *AWS CLI Command Reference*.

For the following examples to work, make sure that you have a subnet in the VPC associated with the route table. The first example describes the route table with the specified VPC ID. The second one describes the route table with the specified route table ID.

```
aws ec2 describe-route-tables \  
--filters "Name=vpc-id,Values=vpc-0daeb67adda59e0cd"  
  
aws ec2 describe-route-tables \  
--route-table-ids rtb-45ac473a
```

The following example describes the specified associations between a specific VPC and your local gateway route tables.

```
aws ec2 describe-local-gateway-route-table-vpc-associations
--filters "Name=vpc-id,Values=vpc-0daeb67adda59e0cd"
```

Amazon Quick elastic network interface

The *Amazon Quick elastic network interface* is a logical networking component in a VPC that represents a virtual network card. Quick creates at least two of these network interfaces to use with a VPC connection based off of the subnets that are attached to it. Then you add the VPC connection to each Amazon Quick Sight data source you create. The Quick network interface alone doesn't give Quick direct access to your databases. The VPC connection works only for the Amazon Quick Sight data sources that are configured to use it.

When you use the Amazon Quick Sight data source to query a database or other instance within your VPC, all the network traffic from Amazon Quick originates from this Amazon Quick network interface. Because the Amazon Quick network interface exists inside your VPC, traffic originating from it can reach destinations within your VPC by using their private IP addresses. Each Amazon Quick network interface gets its own private IP address that comes from the subnet you configure. The private IP address is unique for each AWS account, unlike the public IP range.

Inbound endpoints for Amazon Route 53 Resolver

Amazon Route 53 Resolver provides DNS query capabilities to your VPC. Route 53 Resolver resolves all local DNS queries and recursively looks up any DNS queries that aren't local on public DNS servers.

Amazon Quick can't directly use Route 53 Resolver to query private DNS servers. However, you can set up Route 53 Resolver inbound endpoints to make these queries indirectly. For more information about inbound endpoints, see [Forwarding inbound DNS queries to your VPCs](#) in the *Route 53 Resolver Developer Guide*. To use inbound endpoints in Amazon Quick, provide the IP addresses of the endpoints for **DNS resolver endpoints** when you create a VPC connection.

Finding information to connect to a VPC

Applies to: Enterprise Edition

Intended audience: System administrators

To gather the information to have ready when you create a VPC connection in Amazon Quick Enterprise edition, take the steps listed following.

Steps

- [Identify the data sources to use](#)
- [Identify the AWS Region to use](#)
- [Identify the VPC ID to use](#)
- [Identify the subnet IDs to use](#)
- [Identify the security group to use](#)

Identify the data sources to use

Start by identifying all the data sources that you want to connect to using Quick. For each of these, note the database's private IP, security group, and subnets. Amazon Quick connects to your data using the private IP. However, you don't need to enter this or the security group or subnet information for the VPC connection. This information helps you identify the other components you need for the Amazon Quick VPC connection.

Note

For the connection to your data source to work, make sure that there's a traceable route from your data source to the VPC ID. For more details, see [Identify the data sources to use](#).

Identify the AWS Region to use

For the connection to work, the data, the subnets, and the security group must be in the same VPC. Make sure also that you use Quick in the same AWS Region with the VPC.

You can't use Amazon Quick in one AWS Region and expect to connect to a VPC in a different AWS Region.

If your team is already using Amazon Quick, you can see your current AWS Region displayed at the upper right of the Amazon Quick home screen. You can change the AWS Region you're using in Amazon Quick by changing the Region at the upper right of the Amazon Quick home screen. All the people who plan to use the data in the VPC must be using the same AWS Region in Amazon Quick.

Note

The AWS Region that displays in the Amazon Quick console doesn't have to match your AWS CLI configuration. Take care not to mistake your current Amazon Quick console settings with the settings that apply in any AWS CLI commands that you run or the settings in other consoles. Changing the current AWS Region in any console doesn't change the Region anywhere except for that page.

For example, let's say you have three tabs open in one browser window. You can have the Amazon Quick console open in one AWS Region, the Amazon VPC console open in a second Region, the Amazon RDS console open in a third Region, and the AWS CLI running in a fourth Region.

Identify the VPC ID to use

The VPC ID is assigned when the VPC is created.

Using the AWS CLI

The following `describe-vpcs` example retrieves details for all of your VPCs.

```
aws ec2 describe-vpcs
```

The following `describe-vpcs` example retrieves details for the specified VPC.

```
aws ec2 describe-vpcs \  
--vpc-ids vpc-06e4ab6c6cEXAMPLE
```

Using the Amazon VPC console

In the VPC console (<https://console.aws.amazon.com/vpc/>), choose **Your VPCs** at left. Choose the VPC-ID that you want to use. The correct one has Availability Zones in your AWS Region and also meets the requirements described in [Finding information to connect to a VPC](#). Also note the ID of **Main Route Table**, because you need this to identify related subnets.

Tip

In the Amazon VPC console, you can filter by VPC. This option is located at the top left of the console. If you filter by your VPC ID, all the other menus display only the network elements that are in your selected VPC.

Identify the subnet IDs to use

To locate the subnet IDs for the subnets used by the VPC, open the VPC console. Locate the VPC you are using, and at least two subnets in different availability zones. Amazon Quick creates its Amazon Quick elastic network interface (Amazon Quick network interface) for the subnets that you choose. The Amazon Quick network interfaces get created after you save your VPC connection settings, described in the following section.

Your database instances can reside in different subnets. However, make sure you can trace the route from this subnet to any data destinations that you want to reach.

Using the AWS CLI

The following example describes all existing subnets.

```
aws ec2 describe-subnets
```

The following `describe-subnets` example uses a filter to retrieve details for the subnets of the specified VPC.

```
aws ec2 describe-subnets \
--filters "Name=vpc-id,Values=vpc-06e4ab6c6cEXAMPLE"
```

Using the Amazon VPC console

In the VPC console (<https://console.aws.amazon.com/vpc/>), choose **Subnets** at left, and find the correct **Subnet ID**. Any subnet is correct if your database subnet has a route to the subnet that you choose at this point. In most cases, if you haven't configured the VPC network yourself, all subnets are connected.

Identify the security group to use

The security group contains rules that control the inbound and outbound network traffic on your data source instances. The security group you are using should have the description "QuickSight-VPC" to make it easier to identify.

When you locate the correct security group, copy its **Group ID** value.

Using the AWS CLI

The following example displays the security groups in a specific AWS Region. It displays only the group ID, name, and description. It filters the result to display only groups for a specific VPC ID that also have a description of "QuickSight-VPC".

```
aws ec2 describe-security-groups \
--region us-west-2 \
--query 'SecurityGroups[*].[GroupId, GroupName, Description]' \
--filters "Name=vpc-id,Values=vpc-06e4ab6c6cEXAMPLE"
"Name=description,Values=QuickSight-VPC"
```

The following example displays information about the security group with the ID sg-903004f8. Note that you can't reference a security group for EC2-VPC by name.

```
aws ec2 describe-security-groups
--group-ids sg-903004f8
--region us-west-2
```

The following example queries the results to describe VPC the inbound and outbound rules of a security group with a specific ID (sg-903004f8), in a specific AWS Region (us-west-2).

```
aws ec2 describe-security-groups \
--region us-west-2 \
--group-ids sg-903004f8 \
--query 'SecurityGroups[*].[GroupId, GroupName, Description,
IpPermissions,IpPermissionsEgress]'
```

The following example uses filters to describe VPC security groups that have a specific rule that allows SQL Server traffic (port 1433). The example also has a rule that allows traffic from all addresses (0.0.0.0/0). The output is filtered to display only the group IDs, names, and descriptions of the security groups. Security groups must match all filters to be returned in the results. However, a single rule doesn't have to match all filters. (EC2-VPC only)

```
aws ec2 describe-security-groups \
--filters Name=ip-permission.from-port,Values=1433 \
Name=ip-permission.to-port,Values=1433 \
Name=ip-permission.cidr,Values='0.0.0.0/0' \
--query 'SecurityGroups[*].[GroupId, GroupName, Description]'
```

Using the Amazon VPC console

In the VPC console (<https://console.aws.amazon.com/vpc/>), choose **Security groups** at left, and find the correct group ID. The correct one has your VPC ID on it. It should also have a tag or description that includes the word "QuickSight".

Configuring Amazon Quick Sight access to AWS data sources

Use this section to help you configure access to resources in other AWS services.

We recommend that you use SSL to secure Amazon Quick Sight connections to your data sources. To use SSL, you must have a certificate signed by a recognized certificate authority (CA). Amazon Quick doesn't accept certificates that are self-signed or issued from a nonpublic CA. For more information, see [Amazon Quick SSL and CA certificates](#).

Topics

- [Required permissions](#)
- [Network and database configuration requirements](#)
- [Allowing autodiscovery of AWS resources](#)
- [Authorizing connections from Amazon Quick Sight to AWS data stores](#)
- [Exploring your AWS data in Amazon Quick](#)
- [AWS service connectors](#)

Required permissions

When you connect to a Amazon Quick Sight data source that requires a user name, the user name must have SELECT permissions on some system tables. These permissions allow Amazon Quick Sight to do things such as discover table schemas and estimate table size.

The following table identifies the tables that the account must have SELECT permissions for, depending on the type of database you are connecting to. These requirements apply for all

database instances you connect to, regardless of their environment. In other words, they apply whether your database instances are on-premises, in Amazon RDS, in Amazon EC2, or elsewhere.

Instance type	Tables
Amazon Aurora	INFORMATION_SCHEMA.STATISTICS INFORMATION_SCHEMA.TABLES
Amazon Redshift	pg_stats pg_class pg_namespace
MariaDB	INFORMATION_SCHEMA.STATISTICS INFORMATION_SCHEMA.TABLES
Microsoft SQL Server	DBCC SHOW_STATISTICS sp_statistics
MySQL	INFORMATION_SCHEMA.STATISTICS INFORMATION_SCHEMA.TABLES
Oracle	DBA_TAB_COLS ALL_TABLES dba_segments all_segments user_segments
PostgreSQL	pg_stats pg_class pg_namespace

Instance type	Tables
ServiceNow	<code>sys_dictionary</code> (column metadata) <code>sys_db_object</code> (table metadata) <code>sys_glide_object</code> (field type metadata)

Note

If you are using MySQL or PostgreSQL, verify that you are connecting from an allowed host or IP address. For more detail, see [Database configuration requirements for self-administered instances](#).

Network and database configuration requirements

To serve as data sources, databases need to be configured so that Amazon Quick can access them. Use the following sections to make sure that your database is configured appropriately.

Important

Because a database instance on Amazon EC2 is administered by you rather than AWS, it must meet both the [Network configuration requirements](#) as well as the [Database configuration requirements for self-administered instances](#).

Topics

- [Network configuration requirements](#)
- [Database configuration requirements for self-administered instances](#)

Network configuration requirements

Intended audience: System administrators

For you to use your database server from Amazon Quick, your server must be accessible from the internet. It must also allow inbound traffic from Amazon Quick servers.

If the database is on AWS and in the same AWS Region as your Amazon Quick account, you can auto-discover the instance to make connecting to it easier. To do this, you must grant Amazon Quick permissions to access it. For more information, see [Accessing data sources](#).

Topics

- [Network configuration for an AWS instance in a default VPC](#)
- [Network configuration for an AWS instance in a nondefault VPC](#)
- [Network configuration for an AWS instance in a private VPC](#)
- [Network configuration for an AWS instance that is not in a VPC](#)
- [Network configuration for a database instance other than AWS](#)

Network configuration for an AWS instance in a default VPC

In some cases, your database might be on an AWS cluster or instance that you created in a default VPC. Thus, it's publicly accessible (that is, you didn't choose to make it private). In such cases, your database is already appropriately configured to be accessible from the internet. However, you still need to enable access from Amazon Quick servers to your AWS cluster or instance. For further details on how to do this, choose the appropriate topic following:

- [Authorizing connections from Amazon Quick to Amazon RDS database instances](#)
- [Authorizing connections from Amazon Quick to Amazon Redshift clusters](#)
- [Authorizing connections from Amazon Quick to Amazon EC2 instances](#)

Network configuration for an AWS instance in a nondefault VPC

If you are configuring an AWS instance in a nondefault VPC, make sure that the instance is publicly accessible and that the VPC has the following:

- An internet gateway.
- A public subnet.
- A route in the route table between the internet gateway and the AWS instance.
- Network access control lists (ACLs) in your VPC that allow traffic between the cluster or instance and Amazon Quick servers. These ACLs must do the following:

- Allow inbound traffic from the appropriate Amazon Quick IP address range and all ports to the IP address and port that the database is listening on.
- Allow outbound traffic from the database's IP address and port to the appropriate Amazon Quick IP address range and all ports.

For more information about Amazon Quick IP address ranges, see [IP address ranges for Amazon Quick](#) following.

For more information about configuring VPC ACLs, see [Network ACLs](#).

- Security group rules that allow traffic between the cluster or instance and Amazon Quick servers. For further details on how to create appropriate security group rules, see [Authorizing connections to AWS data sources](#).

For more information about configuring a VPC in the Amazon VPC service, see [Networking in Your VPC](#).

Network configuration for an AWS instance in a private VPC

If your database is on an AWS cluster or instance that you created in a private VPC, you can use it with Amazon Quick. For more information, see [Connecting to a Amazon VPC with Amazon Quick](#).

For more information on Amazon VPC, see [Amazon VPC](#) and [Amazon VPC Documentation](#).

Network configuration for an AWS instance that is not in a VPC

If you are configuring an AWS instance that is not in a VPC, make sure that the instance is publicly accessible. Also, make sure that there is a security group rule that allows traffic between the cluster or instance and Amazon Quick servers. For further details on how to do this, choose the appropriate topic following:

- [Authorizing connections from Amazon Quick to Amazon RDS database instances](#)
- [Authorizing connections from Amazon Quick to Amazon Redshift clusters](#)
- [Authorizing connections from Amazon Quick to Amazon EC2 instances](#)

Network configuration for a database instance other than AWS

To use SSL to secure your connections to your database (*recommended*), make sure that you have a certificate signed by a recognized certificate authority (CA). Amazon Quick doesn't accept

certificates that are self-signed or issued from a nonpublic CA. For more information, see [Amazon Quick SSL and CA certificates](#).

If your database is on a server other than AWS, you must change that server's firewall configuration to accept traffic from the appropriate Amazon Quick IP address range. For more information about Amazon Quick IP address ranges, see [IP address ranges for Amazon Quick](#). For any other steps that you need to take to enable internet connectivity, see your operating system documentation.

Amazon Quick SSL and CA certificates

We recommend that you use a public certificate issued by [AWS Certificate Manager \(ACM\)](#). Amazon Quick supports the same certificate authorities (CAs) as Mozilla, so if you don't use ACM, use a certificate issued by a CA on the [Mozilla Included CA Certificate List](#).

IP address ranges for Amazon Quick

For more information on the IP address ranges for Amazon Quick in supported Regions, see [AWS Regions, websites, IP address ranges, and endpoints](#).

Database configuration requirements for self-administered instances

Intended audience: System administrators and Amazon Quick administrators

For a database to be accessible to Amazon Quick, it must meet the following criteria:

- It must be accessible from the internet. To enable internet connectivity, see your database management system documentation.
- It must be configured to accept connections and authenticate access using the user credentials that you provide as part of creating the data set.
- If you are connecting to MySQL or PostgreSQL, the database engine must be accessible from your host or IP range. This optional security limitation is specified in MySQL or PostgreSQL connection settings. If this limitation is in place, any attempt to connect from a nonspecified host or IP address is rejected, even if you have the correct username and password.
- In MySQL, the server accepts the connection only if the user and host are verified in the user table. For more information, see [Access Control, Stage 1: Connection Verification](#) in the MySQL documentation.

- In PostgreSQL, you control client authentication by using the `pg_hba.conf` file in the database cluster's data directory. However, this file might be named and located differently on your system. For more information, see [Client Authentication](#) in the PostgreSQL documentation.

Allowing autodiscovery of AWS resources

Applies to: Enterprise Edition and Standard Edition

Intended audience: System administrators

Each AWS service that you access from Amazon Quick needs to allow traffic from Quick. Instead of opening each service console separately to add permissions, a Quick administrator can do this in the administration screen. Before you begin, make sure that you have addressed the following prerequisites.

If you choose to enable autodiscovery of AWS resources for your Quick account, Quick creates an AWS Identity and Access Management (IAM) role in your AWS account. This IAM role grants your account permission to identify and retrieve data from your AWS data sources.

Because AWS limits the number of IAM roles that you can create, make sure that you have at least one free role. You need this role for Amazon Quick to use if you want Amazon Quick to autodiscover your AWS resources.

You can have Amazon Quick autodiscover Amazon RDS DB instances or Amazon Redshift clusters that are associated with your AWS account. These resources must be located in the same AWS Region as your Amazon Quick account.

If you choose to enable autodiscovery, choose one of the following options to make the AWS resource accessible:

- For Amazon RDS DB instances that you created in a default VPC and didn't make private, or that aren't in a VPC (EC2-Classic instances), see [Authorizing connections from Amazon Quick to Amazon RDS instances](#). In this topic, you can find information on creating a security group to allow connections from Amazon Quick servers.
- For Amazon Redshift clusters that you created in a default VPC and didn't choose to make private, or that aren't in a VPC (that is, EC2-Classic instances), see [Authorizing connections from](#)

[Amazon Quick to Amazon Redshift clusters](#). In this topic, you can find information on creating a security group to allow connections from Amazon Quick servers.

- For an Amazon RDS DB instance or Amazon Redshift cluster that is in a nondefault VPC, see [Authorizing connections from Amazon Quick to Amazon RDS instances](#) or [Authorizing connections from Amazon Quick to Amazon Redshift clusters](#). In these topics, you can find information on first creating a security group to allow connections from Amazon Quick servers. In addition, you can find information on then verifying that the VPC meets the requirements described in [Network configuration for an AWS instance in a nondefault VPC](#).
- If you don't use a private VPC, set up the Amazon RDS instance to allow connections from the Amazon Quick Region's public IP address.

Enabling autodiscovery is the easiest way to make this data available in Amazon Quick. You can still manually create data connections whether or not you enable autodiscovery.

Authorizing connections from Amazon Quick Sight to AWS data stores

Applies to: Enterprise Edition and Standard Edition

Intended audience: System administrators

For Amazon Quick Sight to access your AWS resources, you must create security groups for them that authorize connections from the IP address ranges used by Amazon Quick Sight servers. You must have AWS credentials that permit you to access these AWS resources to modify their security groups.

Use the procedures in the following sections to enable Amazon Quick Sight connections.

Topics

- [Authorizing connections from Amazon Quick Sight to Amazon RDS DB instances](#)
- [Authorizing connections from Amazon Quick Sight to Amazon Redshift clusters](#)
- [Authorizing connections from Amazon Quick to Amazon EC2 instances](#)
- [Authorizing connections through AWS Lake Formation](#)
- [Authorizing connections to Amazon OpenSearch Service](#)
- [Authorizing connections to Amazon Athena](#)

- [Data access integrations](#)

Authorizing connections from Amazon Quick Sight to Amazon RDS DB instances

Applies to: Enterprise Edition and Standard Edition

Intended audience: System administrators

For Amazon Quick Sight to connect to an Amazon RDS DB instance, you must create a new security group for that DB instance. This security group contains an inbound rule authorizing access from the appropriate IP address range for the Quick servers in that AWS Region. To learn more about authorizing Quick connections, see [Manually enabling access to an Amazon RDS instance in a VPC](#) or [Manually enabling access to an Amazon RDS instance that is not in a VPC](#).

To learn more about authorizing Amazon Quick Sight connections manually, see [Manually enabling access to an Amazon RDS instance in a VPC](#) or [Manually enabling access to an Amazon RDS instance that is not in a Amazon VPC](#).

To create and assign a security group for an Amazon RDS DB instance, you must have AWS credentials that permit access to that DB instance.

Enabling connection from Amazon Quick servers to your instance is just one of several prerequisites for creating a data set based on an AWS database data source. For more information about what is required, see [Creating a dataset from a database](#).

Topics

- [Manually enabling Amazon Quick Sight access to an Amazon RDS instance in a VPC](#)
- [Manually enabling access from Amazon Quick Sight to an Amazon RDS instance that is not in a VPC](#)

Manually enabling Amazon Quick Sight access to an Amazon RDS instance in a VPC

Use the following procedure to enable Amazon Quick Sight access to an Amazon RDS DB instance in a VPC. If your Amazon RDS DB instance is in subnet that is private (in relation to Amazon Quick) or that has Internet Gateways attached, see [Connecting to a VPC with Amazon Quick](#).

To enable Amazon Quick Sight access to an Amazon RDS DB instance in a VPC

1. Sign in to the AWS Management Console and open the Amazon RDS console at <https://console.aws.amazon.com/rds/>.
2. Choose **Databases**, locate the DB instance, and view its details. To do this, you click directly on its name (a hyperlink in the **DB identifier** column).
3. Locate **Port** and note the **Port** value. This can be a number or a range.
4. Locate **VPC** and note the **VPC** value.
5. Choose the **VPC** value to open the VPC console. In the Amazon VPC Management Console, choose **Security Groups** in the navigation pane.
6. Choose **Create Security Group**.
7. On the **Create Security Group** page, enter the security group information as follows:
 - For **Name tag** and **Group name**, enter **Amazon-QuickSight-access**.
 - For **Description**, enter **Amazon-QuickSight-access**.
 - For **VPC**, choose the VPC for your instance. This VPC is the one with the **VPC ID** that you noted previously.
8. Choose **Create**. On the confirmation page, note the **Security Group ID**. Choose **Close** to exit this screen.
9. Choose your new security group from the list, and then choose **Inbound Rules** from the tab list below.
10. Choose **Edit rules** to create a new rule.
11. On the **Edit inbound rules** page, choose **Add rule** to create a new rule.

Use the following values:

- For **Type**, choose **Custom TCP Rule**.
- For **Protocol**, choose **TCP**.
- For **Port Range**, enter the port number or range of the Amazon RDS cluster. This port number (or range) is the one that you noted previously.
- For **Source**, choose **Custom** from the list. Next to the word "Custom", enter the CIDR address block for the AWS Region where you plan to use Amazon Quick.

For example, for Europe (Ireland) you would enter Europe (Ireland)'s CIDR address block: 52.210.255.224/27. For more information on the IP address ranges for Amazon Quick in supported AWS Regions, see [AWS Regions, websites, IP address ranges, and endpoints](#).

 Note

If you have activated Amazon Quick in multiple AWS Regions, you can create inbound rules for each Amazon Quick endpoint CIDR. Doing this allows Amazon Quick to have access to the Amazon RDS DB instance from any AWS Region defined in the inbound rules.

Anyone who uses Amazon Quick in multiple AWS Regions is treated as a single user. In other words, even if you are using Amazon Quick in every AWS Region, both your Amazon Quick subscription (sometimes called an 'account') and your users are global.

12. For **Description**, enter a useful description, for example "*Europe (Ireland) QuickSight*".
13. Choose **Save rules** to save your new inbound rule. Then choose **Close**.
14. Go back to the detailed view of the DB instance. Return the Amazon RDS console (<https://console.aws.amazon.com/rds/>) and choose **Databases**.
15. Choose the DB identifier for the relevant RDS instance. Choose **Modify**. The same screen displays whether you choose Modify from the databases screen or the DB instance screen: **Modify DB Instance**.
16. Locate the **Network & Security** section (the third section from the top).

The currently assigned security group or groups are already chosen for **Security Group**. Don't remove any of the existing ones unless you are sure.

Instead, choose your new security group to add it to the other groups that are selected. If you followed the name suggested previously, this group might be named something similar to **Amazon-QuickSight-access**.

17. Scroll to the bottom of the screen. Choose **Continue**. and then choose **Modify DB Instance**.
18. Choose **Apply during the next scheduled maintenance** (the screen indicates when this will occur).

Don't choose **Apply immediately**. Doing this also applies any additional changes that are in the pending modifications queue. Some of these changes might require downtime. If you bring

the server down outside the maintenance window, this can cause a problem for users of this DB instance. Consult your system administrators before applying immediate changes.

19. Choose **Modify DB Instance** to confirm your changes. Then, wait for the next maintenance window to pass.

Manually enabling access from Amazon Quick Sight to an Amazon RDS instance that is not in a VPC

Use the following procedure to access an Amazon RDS DB instance that is not in a VPC. You can associate a security group with a DB instance by using **Modify** on the RDS console, the `ModifyDBInstance` Amazon RDS API, or the `modify-db-instance` AWS CLI command.

Note

This section included for backwards compatibility purposes.

To use the console to access an Amazon RDS DB instance that is not in a VPC

1. Sign in to the AWS Management Console and open the Amazon RDS console at <https://console.aws.amazon.com/rds/>.
2. Choose **Databases**, select the DB instance, and choose **Modify**.
3. Choose **Security Groups** in the navigation pane.
4. Choose **Create DB Security Group**.
5. Enter **Amazon-QuickSight-access** for the **Name** and **Description** values, and then choose **Create**.
6. The new security group is selected by default.

Select the details icon next to the security group, as shown following.

7. For **Connection Type**, choose **CIDR/IP**.
8. For **CIDR/IP to Authorize**, enter the appropriate CIDR address block. For more information on the IP address ranges for Amazon Quick in supported AWS Regions, see [AWS Regions, websites, IP address ranges, and endpoints](#).
9. Choose **Authorize**.
10. Return to the **Instances** page of the Amazon RDS Management Console, choose the instance that you want to enable access to, choose **Instance Actions**, and then choose **Modify**.

11. In the **Network & Security** section, the currently assigned security group or groups already is chosen for **Security Group**. Press CTRL and choose **Amazon-QuickSight-access** in addition to the other selected groups.
12. Choose **Continue**, and then choose **Modify DB Instance**.

Authorizing connections from Amazon Quick Sight to Amazon Redshift clusters

Applies to: Enterprise Edition and Standard Edition

Intended audience: System administrators

You can provide access to Amazon Redshift data using three authentication methods: trusted identity propagation, run-as IAM role, or Amazon Redshift database credentials.

With trusted identity propagation, a user's identity is passed to Amazon Redshift with single sign-on that is managed by IAM Identity Center. A user that accesses a dashboard in Amazon Quick Sight has their identity propagated to Amazon Redshift. In Amazon Redshift, fine grained data permissions are applied on the data before the data is presented in a Amazon Quick asset to the user. Amazon Quick authors can also connect to Amazon Redshift data sources without a password input or IAM role. If Amazon Redshift Spectrum is used, all permission management is centralized in Amazon Redshift. Trusted identity propagation is supported when Amazon Quick and Amazon Redshift use the same organization instance of IAM Identity Center. Trusted identity propagation is not currently supported for the following features.

- SPICE datasets
- Custom SQL on data sources
- Alerts
- Email reports
- Amazon Quick Q
- CSV, Excel, and PDF exports
- Anomaly detection

For Amazon Quick to connect to an Amazon Redshift instance, you must create a new security group for that instance. This security group contains an inbound rule that authorizes access from the appropriate IP address range for the Amazon Quick servers in that AWS Region. To learn more about authorizing Amazon Quick connections, see [Manually enabling access to an Amazon Redshift cluster in a VPC](#).

Enabling connection from Amazon Quick servers to your cluster is just one of several prerequisites for creating a data set based on an AWS database data source. For more information about what is required, see [Creating a dataset from a database](#).

Topics

- [Enabling trusted identity propagation with Amazon Redshift](#)
- [Manually enabling access to an Amazon Redshift cluster in a VPC](#)
- [Enabling access to Amazon Redshift Spectrum](#)

Enabling trusted identity propagation with Amazon Redshift

Trusted identity propagation authenticates the end user in Amazon Redshift when they access Amazon Quick assets that leverage a trusted identity propagation enabled data source. When an author creates a data source with trusted identity propagation, the identity of the data source consumers in Amazon Quick Sight is propagated and logged in CloudTrail. This allows database administrators to centrally manage data security in Amazon Redshift and automatically apply all data security rules to data consumers in Amazon Quick. With other authentication methods, the data permissions of the author who created the data source are applied to all data source consumers. The data source author can choose to apply additional row and column level security to the data sources that they create in Amazon Quick Sight.

Trusted identity propagation data sources are supported only in Direct Query datasets. SPICE datasets do not currently support trusted identity propagation.

Topics

- [Prerequisites](#)
- [Enabling trusted identity propagation in Amazon Quick Sight](#)
- [Connecting to Amazon Redshift with trusted identity propagation](#)

Prerequisites

Before you get started, make sure that you have all of the required prerequisites ready.

- Trusted identity propagation is only supported for Amazon Quick accounts that are integrated with IAM Identity Center. For more information, see [Configure your Amazon Quick account with IAM Identity Center](#).
- An Amazon Redshift application that is integrated with IAM Identity Center. The Amazon Redshift cluster that you use must be in the same organization in AWS Organizations as the Amazon Quick account that you want to use. The cluster must also be configured with the same organization instance in IAM Identity Center that your Amazon Quick account is configured to. For more information about configuring a Amazon Redshift cluster, see [Integrating IAM Identity Center](#).

Enabling trusted identity propagation in Amazon Quick Sight

To configure Amazon Quick Sight to connect to Amazon Redshift data sources with trusted identity propagation, configure Amazon Redshift OAuth scopes to your Amazon Quick account.

To add a scope that allows Amazon Quick to authorize identity propagation to Amazon Redshift, specify the AWS account ID of the Amazon Quick account and the service that you want to authorize identity propagation with, in this case 'REDSHIFT'.

Specify the IAM Identity Center application ARN of the Amazon Redshift cluster that you are authorizing Amazon Quick to propagate user identities to. This information can be found in the Amazon Redshift console. If you don't specify authorized targets for the Amazon Redshift scope, Amazon Quick authorizes users from any Amazon Redshift cluster that share the same IAM Identity Center instance. The example below configures Amazon Quick to connect to Amazon Redshift data sources with trusted identity propagation.

```
aws quicksight update-identity-propagation-config --aws-account-id "AWSACCOUNTID"  
--service "REDSHIFT" --authorized-targets "arn:aws:sso::XXXXXXXXXXXX:application/  
ssoins-XXXXXXXXXXXX/ap1-XXXXXXXXXXXX" "arn:aws:sso::XXXXXXXXXXXX:application/  
ssoins-XXXXXXXXXXXX/ap1-XXXXXXXXXXXX"
```

The following example deletes OAuth scopes from a Amazon Quick account.

```
aws quicksight delete-identity-propagation-config --aws-account-id "AWSACCOUNTID"  
--service "REDSHIFT" --authorized-targets "arn:aws:sso::XXXXXXXXXXXX:application/
```

```
ssoins-XXXXXXXXXXXXap1-XXXXXXXXXXXX "arn:aws:sso::XXXXXXXXXXXX:application/  
ssoins-XXXXXXXXXXXX/ap1-XXXXXXXXXXXX"
```

The following example lists all OAuth scopes that are currently on a Amazon Quick account.

```
aws quicksight list-identity-propagation-configs --aws-account-id "AWSACCOUNTID"
```

Connecting to Amazon Redshift with trusted identity propagation

Use the procedure below to connect to Amazon Redshift trusted identity propagation.

To connect to Amazon Redshift with trusted identity propagation

1. Create a new dataset in Amazon Quick. For more information about creating a dataset, see [Creating datasets](#).
2. Choose Amazon Redshift as the data source for the new dataset.

Note

The authentication type of an existing data source can't be changed to trusted identity propagation

3. Choose IAM Identity Center as the identity option for the data source, and then choose **Create data source**.

Manually enabling access to an Amazon Redshift cluster in a VPC

Applies to: Enterprise Edition

Use the following procedure to enable Amazon Quick Sight access to an Amazon Redshift cluster in a VPC.

To enable Amazon Quick Sight access to an Amazon Redshift cluster in a VPC

1. Sign in to the AWS Management Console and open the Amazon Redshift console at <https://console.aws.amazon.com/redshiftv2/>.
2. Navigate to the cluster that you want to make available in Amazon Quick.
3. In the **Cluster Properties** section, find **Port**. Note the **Port** value.

4. In the **Cluster Properties** section, find **VPC ID** and note the **VPC ID** value. Choose **VPC ID** to open the Amazon VPC console.
5. On the Amazon VPC console, choose **Security Groups** in the navigation pane.
6. Choose **Create Security Group**.
7. On the **Create Security Group** page, enter the security group information as follows:
 - For **Security group name**, enter **redshift-security-group**.
 - For **Description**, enter **redshift-security-group**.
 - For **VPC**, choose the VPC for your Amazon Redshift cluster. This is the VPC with the VPC ID that you noted.
8. Choose **Create security group**.

Your new security group should appear on the screen.

9. Create a second security group with the following properties.
 - For **Security group name**, enter **quicksight-security-group**.
 - For **Description**, enter **quicksight-security-group**.
 - For **VPC**, choose the VPC for your Amazon Redshift cluster. This is the VPC with the VPC ID that you noted.
10. Choose **Create security group**.
11. After you create the new security groups, create inbound rules for the new groups.

Choose the new **redshift-security-group** security group, and input the following values.

- For **Type**, choose **Amazon Redshift**.
 - For **Protocol**, choose **TCP**.
 - For **Port Range**, enter the port number of the Amazon Redshift cluster to which you are providing access. This is the port number that you noted in an earlier step.
 - For **Source**, enter the security group ID of **quicksight-security-group**.
12. Choose **Save rules** to save your new inbound rule.
 13. Repeat the previous step for **quicksight-security-group** and enter the following values.
 - For **Type**, choose **All traffic**.
 - For **Protocol**, choose **All**.
 - For **Port Range**, choose **All**.

- For **Source**, enter the security group ID of `redshift-security-group`.
14. Choose **Save rules** to save your new inbound rule.
 15. In Amazon Quick, navigate to the **Manage Amazon Quick** menu.
 16. Choose **Manage VPC connections**, and then choose **Add VPC connection**.
 17. Configure the new VPC connection with the following values.
 - For **VPC connection name**, choose a meaningful name for the VPC connection.
 - For **VPC ID**, choose the VPC in which the Amazon Redshift cluster exists.
 - For **Subnet ID**, choose the subnet for the Availability Zone (AZ) that is used for Amazon Redshift.
 - For **Security group id**, copy and paste the security group ID for `quicksight-security-group`.
 18. Choose **Create**. It might take several minutes for the new VPC to generate.
 19. In the Amazon Redshift console, navigate to the Amazon Redshift cluster that `redshift-security-group` is configured to. Choose **Properties**. under **Network and security settings**, enter the name of the security group.
 20. In Amazon Quick, choose **Datasets**, and then choose **New dataset**. Create a new dataset with the following values.
 - For **Data source**, choose **Amazon Redshift Auto-discovered**.
 - Give the data source a meaningful name.
 - The instance ID should auto populate with the VPC connection that you created in Amazon Quick. If the instance ID doesn't auto populate, choose the VPC that you created from the dropdown list.
 - Enter the database credentials. If your Amazon Quick account uses trusted identity propagation, choose **Single sign-on**.
 21. Validate the connection, and then choose **Create data source**.

If you want to restrict the default outbound rules further, update the outbound rule of `quicksight-security-group` to allow only Amazon Redshift traffic to `redshift-security-group`. You can also delete the outbound rule that's located in the `redshift-security-group`.

Enabling access to Amazon Redshift Spectrum

Using Amazon Redshift Spectrum, you can connect Amazon Quick to an external catalog with Amazon Redshift. For example, you can access the Amazon Athena catalog . You can then query unstructured data on your Amazon S3 data lake using an Amazon Redshift cluster instead of the Athena query engine.

You can also combine data sets that include data stored in Amazon Redshift and in S3. Then you can access them using the SQL syntax in Amazon Redshift.

After you've registered your data catalog (for Athena) or external schema (for a [Hive metastore](#)), you can use Amazon Quick to choose the external schema and Amazon Redshift Spectrum tables. This process works just as for any other Amazon Redshift tables in your cluster. You don't need to load or transform your data.

For more information on using Amazon Redshift Spectrum, see [Using Amazon Redshift Spectrum to query external data](#) in the *Amazon Redshift Database Developer Guide*.

To connect using Redshift Spectrum, do the following:

- Create or identify an IAM role associated with the Amazon Redshift cluster.
- Add the IAM policies `AmazonS3ReadOnlyAccess` and `AmazonAthenaFullAccess` to the IAM role.
- Register an external schema or data catalog for the tables that you plan to use.

Redshift Spectrum lets you separate storage from compute, so you can scale them separately. You only pay for the queries that you run.

To connect to Redshift Spectrum tables, you don't need to grant Amazon Quick access to Amazon S3 or Athena. Amazon Quick needs access only to the Amazon Redshift cluster. For full details on configuring Redshift Spectrum, see [Getting started with Amazon Redshift Spectrum](#) in the *Amazon Redshift Database Developer Guide*.

Authorizing connections from Amazon Quick to Amazon EC2 instances

Applies to: Enterprise Edition and Standard Edition

Intended audience: System administrators

For Amazon Quick Sight to connect to an Amazon EC2 instance, you must create a new security group for that instance. This security group contains an inbound rule authorizing access from the appropriate IP address range for the Quick servers in that AWS Region.

To modify the security groups for these Amazon EC2 instances, you must have AWS credentials that permit you to access to the instances.

Enabling connection from Quick servers to your instance is just one of several prerequisites for creating a data set based on an AWS database data source. For more information about what is required, see [Creating a dataset from a database](#).

To enable Amazon Quick access to an Amazon EC2 instance

1. Sign in to the AWS Management Console and open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. If your EC2 instance is in a VPC, choose the instance to view the instance details pane. Find its VPC ID and note that ID for later use.
3. Choose **Security Groups** in the **NETWORK & SECURITY** section of the navigation pane. Then choose **Create Security Group**, as shown following.
4. Enter the security group information as follows:
 - For **Security group name**, enter **Amazon-QuickSight-access**.
 - For **Description**, enter **Amazon-QuickSight-access**.
 - For **VPC**, choose the VPC ID that you noted in step 2 if your Amazon EC2 instance is in a VPC. Otherwise, choose **No VPC**.
5. Choose **Add Rule** on the **Inbound** tab.
6. Create a new rule with the following values:
 - For **Type**, choose **Custom TCP Rule**.
 - For **Protocol**, choose **TCP**.
 - (Optional) For **Port Range**, enter the port number used by the instance on this Amazon EC2 instance to which you are providing access.
 - For **Source**, enter the CIDR address block for the AWS Region where you plan to use Amazon Quick. For example, here is the CIDR address block for Europe (Ireland): 52.210.255.224/27. For more information on the IP address ranges for Amazon Quick in supported AWS Regions, see [AWS Regions, websites, IP address ranges, and endpoints](#).

Note

If you have activated Amazon Quick in multiple AWS Regions, you can create inbound rules for each Amazon Quick endpoint CIDR. Doing this allows Amazon Quick to have access to the Amazon RDS DB instance from any AWS Region defined in the inbound rules.

An Amazon Quick user or administrator who uses Amazon Quick in multiple AWS Regions is treated as a single user. In other words, even if you are using Amazon Quick in every AWS Region, both your Amazon Quick account and your users are global.

7. Choose **Create**.
8. Choose **Instances** in the **INSTANCES** section of the navigation pane, and then choose the instance that you want to enable access to.
9. Choose **Actions**, then **Networking**, and then **Change Security Groups**.
10. In **Change Security Groups**, choose the **Amazon-QuickSight-access** security group.

Then choose **Assign Security Groups**, as shown following.

Authorizing connections through AWS Lake Formation

Applies to: Enterprise Edition


Intended audience: System administrators

If you are querying data with Amazon Athena, you can use AWS Lake Formation to simplify how you secure and connect to your data from Amazon Quick Sight. Lake Formation adds to the AWS Identity and Access Management (IAM) permissions model by providing its own permissions model that is applied to AWS analytics and machine learning services. This centrally defined permissions model controls data access at a granular level through a simple grant and revoke mechanism. You can use Lake Formation instead of, or in addition to, using scoped-down policies with IAM.

When you set up Lake Formation, you register your data sources to allow it to move the data into a new data lake in Amazon S3. Lake Formation and Athena both work seamlessly with AWS Glue

Data Catalog, making it easy to use them together. Athena databases and tables are metadata containers. These containers describe the underlying schema of the data, the data definition language (DDL) statements, and the location of the data in Amazon S3.

The following diagram shows the relationships of the AWS services involved.



After Lake Formation is configured, you can use Amazon Quick to access databases and tables by name or through SQL queries. Amazon Quick provides a full-featured editor where you can write SQL queries. Or you can use the Athena console, the AWS CLI, or your favorite query editor. For more information, see [Accessing Athena](#) in the *Amazon Athena User Guide*.

Use the topics below to configure a Lake Formation connection through Lake Formation or through Amazon Quick.

Topics

- [Enabling connection from Lake Formation](#)
- [Enabling connection from Amazon Quick](#)

Enabling connection from Lake Formation

Before you begin using this solution with Quick, make sure that you can access your data using Athena with Lake Formation. After you verify that the connection is working through Athena, you need to verify only that Amazon Quick can connect to Athena. Doing this means you don't have to troubleshoot connections through all three products at once. One easy way to test the connection is to use the [Athena query console](#) to run a simple SQL command, for example `SELECT 1 FROM table`.

To set up Lake Formation, the person or team who works on it needs access to create a new IAM role and to Lake Formation. They also need the information shown in the following list. For more information, see [Setting up lake formation](#) in the *AWS Lake Formation Developer Guide*.

- Collect the Amazon Resource Names (ARNs) of the Quick users and groups that need to access the data in Lake Formation. These users should be Amazon Quick authors or administrators.

To find Amazon Quick user and group ARNs

1. Use the AWS CLI to find user ARNs for Amazon Quick authors and admins. To do this, run the following `list-users` command in your terminal (Linux or Mac) or at your command prompt (Windows).

```
aws quicksight list-users --aws-account-id 111122223333 --namespace default --  
region us-east-1
```

The response returns information for each user. We show the Amazon Resource Name (ARN) in bold in the following example.

```
RequestId: a27a4cef-4716-48c8-8d34-7d3196e76468  
Status: 200  
UserList:  
- Active: true
```

```

Arn: arn:aws:quicksight:us-east-1:111122223333:user/default/SaanviSarkar
Email: SaanviSarkar@example.com
PrincipalId: federated/iam/AIDAJVCZ0VSR3DESMJ7TA
Role: ADMIN
UserName: SaanviSarkar

```

To avoid using the AWS CLI, you can construct the ARNs for each user manually.

2. (Optional) Use the AWS CLI to find ARNs for Amazon Quick groups by running the following `list-group` command in your terminal (Linux or Mac) or at your command prompt (Windows).

```
aws quicksight list-groups --aws-account-id 111122223333 --namespace default --region us-east-1
```

The response returns information for each group. The ARN appears in bold in the following example.

```

GroupList:
- Arn: arn:aws:quicksight:us-east-1:111122223333:group/default/DataLake-Scorecard
  Description: Data Lake for CX0 Balanced Scorecard
  GroupName: DataLake-Scorecard
  PrincipalId: group/d-90671c9c12/6f9083c2-8400-4389-8477-97ef05e3f7db
  RequestId: c1000198-18fa-4277-a1e2-02163288caf6
  Status: 200

```

If you don't have any Amazon Quick groups, add a group by using the AWS CLI to run the `create-group` command. There currently isn't an option to do this from the Amazon Quick console. For more information, see [Creating and managing groups in Amazon Quick](#).

To avoid using the AWS CLI, you can construct the ARNs for each group manually.

Enabling connection from Amazon Quick

To work with Lake Formation and Athena, make sure that you have AWS resource permissions configured in Amazon Quick:

- Enable access to Amazon Athena.

- Enable access to the correct buckets in Amazon S3. Usually S3 access is enabled when you enable Athena. However, because you can change S3 permissions outside of that process, it's a good idea to verify them separately.

For information about how to verify or change AWS resource permissions in Quick, see [Allowing autodiscovery of AWS resources](#) and [Accessing data sources](#).

Authorizing connections to Amazon OpenSearch Service

Applies to: Enterprise Edition

Intended audience: System administrators

Before you can use OpenSearch in a Amazon Quick Sight dataset, there are a few tasks for the Quick administrator to complete with the cooperation of a person who has access to the OpenSearch console.

To get started, identify each OpenSearch domain that you want to connect to. Then gather the following information for each domain:

- The name of the OpenSearch domain.
- The OpenSearch version used by this domain.
- The Amazon Resource Name (ARN) of the OpenSearch domain.
- The HTTPS endpoint.
- The OpenSearch Dashboards URL, if you use Dashboards. You can extrapolate the Dashboards URL by appending `"/dashboards/"` to an endpoint.
- If the domain has a VPC endpoint, gather all the related information on the VPC tab of the OpenSearch Service console:
 - The VPC ID
 - The VPC security groups
 - The associated IAM role or roles
 - The associated Availability Zones
 - The associated subnets

- If the domain has a regular endpoint (not a VPC endpoint), note that it uses the public network.
- The start hour for the daily automated snapshot (if your users want to know).

Before you proceed, the Amazon Quick administrator enables authorized connections from Amazon Quick to OpenSearch Service. This process is required for every AWS service that you connect to from Amazon Quick. You need to do this only once per AWS account for each AWS service that you use as a data source.

For OpenSearch Service, the authorization process adds the AWS managed policy `AWSQuickSightOpenSearchPolicy` to your AWS account.

Important

Make sure that the IAM policy for your OpenSearch domain doesn't conflict with the permissions in `AWSQuickSightOpenSearchPolicy`. You can find the domain access policy in the OpenSearch Service console. For more information, see [Configuring access policies](#) in the *Amazon OpenSearch Service Developer Guide*.

To turn on or turn off connections from Amazon Quick to OpenSearch Service

1. Within Amazon Quick, choose **Administrator** and **Manage Amazon Quick**.
2. Choose **Security & permissions, Add or remove**.
3. To enable connections, select the **Amazon OpenSearch Service** check box.

To disable connections, clear the **Amazon OpenSearch Service** check box.

4. Choose **Update** to confirm your choices.

If needed, use the topics below to configure a OpenSearch VPC connection and permissions for Amazon Quick to access OpenSearch.

Topics

- [Using a VPC connection](#)
- [Using OpenSearch permissions](#)

Using a VPC connection

In some cases, your OpenSearch domain is in a virtual private cloud (VPC) based on the Amazon VPC service. If so, make sure to determine if Amazon Quick is already connected to the VPC ID that the OpenSearch domain uses. You can reuse an existing VPC connection. If you're not sure if it's working, you can test it. For more information, see [Testing the connection to your Amazon VPC data source](#).

If a connection isn't already defined in Amazon Quick for the VPC that you want to use, you can create one. This task is a multistep process that you need to complete before you proceed. To learn how to add Amazon Quick to a VPC and add a connection from Amazon Quick to the VPC, see [Connecting to a Amazon VPC with Amazon Quick](#).

Using OpenSearch permissions

After you configure Amazon Quick to connect to OpenSearch Service, you might need to enable permissions in OpenSearch. For this part of the setup process, you can use the OpenSearch Dashboards link for each OpenSearch domain. Use the following list to help determine what permissions you need:

1. For domains that use fine-grained access control, configure permissions in the form of a role. This process is similar to using scoped-down policies in Amazon Quick.
2. For each domain that you create a role for, add a role mapping.

For more information, see following.

If your OpenSearch domain has [fine-grained access control](#) enabled, there are some permissions to configure so the domain is accessible from Amazon Quick. Perform these steps for each domain that you want to use.

The following procedure uses OpenSearch Dashboards, which is an open-source tool that works with OpenSearch. You can find the link to Dashboards on the domain dashboard on the OpenSearch Service console.

To add permissions to a domain to allow access from Amazon Quick

1. Open OpenSearch Dashboards for the OpenSearch domain that you want to work with. The URL is `opensearch-domain-endpoint/dashboards/`.
2. Choose **Security** from the navigation pane.

If you don't see the navigation pane, open it by using the menu icon at upper left. To keep the menu open, choose **Dock navigation** at lower left.

3. Choose **Roles, Create role**.
4. Name the role **quicksight_role**.

You can choose a different name, but we recommend this one because we use it in our documentation and it's thus easier to support.

5. Under **Cluster permissions**, add the following permissions:

- `cluster:monitor/main`
- `cluster:monitor/health`
- `cluster:monitor/state`
- `indices:data/read/scroll`
- `indices:data/read/scroll/clear,`

6. Under **Index permissions** specify `*` as the index pattern.

7. For **Index permissions**, add the following permissions:

- `indices:admin/get`
- `indices:admin/mappings/get`
- `indices:admin/mappings/fields/get*`
- `indices:data/read/search*`
- `indices:monitor/settings/get`

8. Choose **Create**.
9. Repeat this procedure for each OpenSearch domain that you're planning to use.

Use the following procedure to add a role mapping for the permissions that you added in the previous procedure. You might find it more efficient to add the permissions and the role mapping as part of a single process. These instructions are separate for clarity.

To create a role mapping for the IAM role you added

1. Open OpenSearch Dashboards for the OpenSearch domain that you want to work with. The URL is *opensearch-domain-endpoint*/dashboards/.
2. Choose **Security** from the navigation pane.

3. Search for and open **quicksight_role** from the list.
4. On the **Mapped users** tab, choose **Manage mapping**.
5. In the **Backend roles** section, enter the ARN of the AWS-managed IAM role for Amazon Quick. Following is an example.

```
arn:aws:iam::AWS-ACCOUNT-ID:role/service-role/aws-quicksight-service-role-v0
```

6. Choose **Map**.
7. Repeat this procedure for each OpenSearch domain that you want to use.

Authorizing connections to Amazon Athena

If you need use Amazon Quick Sight with Amazon Athena or Amazon Athena Federated Query, you first need to authorize connections to Athena and the associated buckets in Amazon Simple Storage Service (Amazon S3). Amazon Athena is an interactive query service that makes it easy to analyze data directly in Amazon S3 using standard SQL. Athena Federated Query provides access to more types of data by using AWS Lambda. Using a connection from Quick to Athena, you can write SQL queries to interrogate data that's stored in relational, non-relational, object, and custom data sources. For more information, see [Using Athena federated query](#) in the Amazon Athena User Guide.

Review the following considerations when setting up access to Athena from Quick:

- Athena stores query results from Amazon Quick Sight in a bucket. By default, this bucket has a name similar to `aws-athena-query-results-AWSREGION-AWSACCOUNTID`, for example `aws-athena-query-results-us-east-2-111111111111`. Therefore, it's important to make sure Amazon Quick Sight has permissions to access the bucket Athena is currently using.
- If your data file is encrypted with an AWS KMS key, grant permissions to the Amazon Quick Sight IAM role to decrypt the key. The easiest way to do this is to use the AWS CLI.

You can run the KMS [create-grant](#) API operation in AWS CLI to do this.

```
aws kms create-grant --key-id <KMS_KEY_ARN> /  
--grantee-principal <QS_ROLE_ARN> --operations Decrypt
```

The Amazon Resource Name (ARN) for the Amazon Quick role has the format
`arn:aws:iam::<account id>:role/service-role/aws-quicksight-s3-consumers-`

role-v<version number> and can be accessed from the IAM console. To find your KMS key ARN, use the S3 console. Go to the bucket that contains your data file and choose the **Overview** tab. The key is located near **KMS key ID**.

- For Amazon Athena, Amazon S3, and Athena Query Federation connections, Amazon Quick uses the following IAM role by default:

```
arn:aws:iam::AWS-ACCOUNT-ID:role/service-role/aws-quicksight-s3-consumers-role-v0
```

If the aws-quicksight-s3-consumers-role-v0 is not present, then Amazon Quick uses:

```
arn:aws:iam::AWS-ACCOUNT-ID:role/service-role/aws-quicksight-service-role-v0
```

- If you assigned scope-down policies to your users, verify that the policies contain the `lambda:InvokeFunction` permission. Without this permission, your users can't access Athena Federated Queries. For more information about assigning IAM policies to your users in Amazon Quick, see [Setting granular access to AWS services through IAM](#). For more information about the `lambda:InvokeFunction` permission, see [Actions, resources, and condition keys for AWS Lambda](#) in the *IAM User Guide*.

To authorize Amazon Quick to connect to Athena or Athena federated data sources

1. (Optional) If you are using AWS Lake Formation with Athena, you also need to enable Lake Formation. For more information, see [Authorizing connections through AWS Lake Formation](#).
2. Open your profile menu at top right and choose **Manage QuickSight**. You must be a Amazon Quick administrator to do this. If you don't see **Manage QuickSight** on the profile menu, you don't have sufficient permissions.
3. Choose **Security & permissions, Add or remove**.
4. Choose the box near Amazon Athena, **Next**.

If it was already enabled, you might have to double-click it. Do this even if Amazon Athena is already enabled, so you can view the settings. No changes are saved until you choose **Update** at the end of this procedure.

5. Enable the S3 buckets you want to access.
6. (Optional) To enable Athena federated queries, select the Lambda functions you want to use.

Note

You can only see Lambda functions for the Athena catalogs in the same region of Amazon Quick.

7. To confirm your changes, choose **Finish**.

To cancel, choose **Cancel**.

8. To save changes to security and permissions, choose **Update**.

To test the connection authorization settings

1. From the Amazon Quick start page, choose **Datasets, New dataset**.
2. Choose the Athena card.
3. Follow the screen prompts to create a new Athena data source using the resources you need to connect to. Choose **Validate connection** to test the connection.
4. If the connection validates, you have successfully configured an Athena or Athena Federated Query connection.

If you don't have sufficient permissions to connect to an Athena dataset or run an Athena query, an error displays directing you to contact a Amazon Quick administrator. This error means need to recheck your connection authorization settings to find the discrepancy.

5. After you can connect successfully, you or your Amazon Quick authors can create data sources connections and share them with other Amazon Quick authors. The authors can then create multiple datasets from the connections, to use in Amazon Quick dashboards.

For troubleshooting information on Athena, see [Connectivity issues when using Athena with Amazon Quick](#).

Using trusted identity propagation with Athena

Trusted identity propagation gives AWS services access to AWS resources based on the user's identity context and securely shares this user's identity with other AWS services. These capabilities enable user access to be more easily defined, granted, and logged.

When administrators configure Quick, Athena, Amazon S3 Access Grants, and AWS Lake Formation with IAM Identity Center, they can now enable trusted identity propagation across these services

and allow the user's identity to be propagated across services. When data is accessed from Quick by an IAM Identity Center user, Athena or Lake Formation can make authorization decisions using the permissions defined for their user or group membership from the organization's identity provider.

Trusted identity propagation with Athena only works when permissions are managed through Lake Formation. User permissions to data reside in Lake Formation.

Prerequisites

Before you get started, make sure that you have the following required prerequisites completed.

Important

As you complete the following prerequisites, note that your IAM Identity Center instance, Athena workgroup, Lake Formation and Amazon S3 Access Grants must all be deployed in the same AWS Region.

- Configure your Quick account with IAM Identity Center. Trusted identity propagation is only supported for Quick accounts that are integrated with IAM Identity Center. For more information, see [Configure your Amazon Quick account with IAM Identity Center](#).

Note

To create Athena data sources, you must be an IAM Identity Center user (author) in a Quick account that uses IAM Identity Center.

- An Athena workgroup that is enabled with IAM Identity Center. The Athena workgroup that you use must be using the same IAM Identity Center instance as the Quick account. For more information about configuring an Athena workgroup, see [Creating an IAM Identity Center enabled Athena workgroup](#) in the *Amazon Athena User Guide*.
- Access to Athena query results bucket is managed with Amazon S3 Access Grants. For more details, see [Managing access with Amazon S3 Access Grants](#) in the *Amazon S3 User Guide*. If your query results are encrypted with an AWS KMS key, the Amazon S3 Access Grant IAM role and the Athena workgroup role both need permissions to AWS KMS.
- For more information, see [Amazon S3 Access Grants and corporate directory identities](#) in the *Amazon S3 User Guide*.

- The Amazon S3 Access Grant role should have the STS: SetContext action in its trust policy for identity propagation. To see an example, see [Register a location](#) in the Amazon S3 User Guide.
- Permissions to data must be managed with Lake Formation and Lake Formation must be configured with the same IAM Identity Center instance as Quick and the Athena workgroup. For configuration information, see [Integrating IAM Identity Center](#) in the *AWS Lake Formation Developer Guide*.
- The data lake administrator needs to grant permissions to IAM Identity Center users and groups in Lake Formation. For more details, [Granting permissions to users and groups](#) in the *AWS Lake Formation Developer Guide*.
- The Quick administrator needs to authorize connections to Athena. For details, see [Authorizing connections to Amazon Athena](#). Note, with trusted identity propagation, you do not need to give the Quick role Amazon S3 bucket permissions or AWS KMS permissions. You need to keep your users and groups that have permissions to the workgroup in Athena in sync with the Amazon S3 bucket that stores query results with Amazon S3 Access Grants permissions so that users can successfully run queries and retrieve query results in the Amazon S3 bucket using trusted identity propagation.

Configure IAM role with required permissions

To use trusted identity propagation with Athena, your Quick account must have the required permissions to access your resources. To provide those permissions, you must configure your Quick account to use an IAM role with the permissions.

If your Quick account is already using a custom IAM role, you can modify that one. If you do not have an existing IAM role, create one by following the instructions in [Create a role for an IAM user](#) in the *IAM User Guide*.

The IAM role you create or modify must contain the following trust policy and permissions.

Required trust policy

For information about updating the trust policy of an IAM role, see [Update a role trust policy](#).

Required Athena permissions

For information about updating the trust policy of an IAM role, see [Update permissions for a role](#).

Note

The Resource uses the * wildcard. We recommend that you update it to include only the Athena resources you want to use with Quick.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "athena:BatchGetQueryExecution",
        "athena:CancelQueryExecution",
        "athena:GetCatalogs",
        "athena:GetExecutionEngine",
        "athena:GetExecutionEngines",
        "athena:GetNamespace",
        "athena:GetNamespaces",
        "athena:GetQueryExecution",
        "athena:GetQueryExecutions",
        "athena:GetQueryResults",
        "athena:GetQueryResultsStream",
        "athena:GetTable",
        "athena:GetTables",
        "athena:ListQueryExecutions",
        "athena:RunQuery",
        "athena:StartQueryExecution",
        "athena:StopQueryExecution",
        "athena:ListWorkGroups",
        "athena:ListEngineVersions",
        "athena:GetWorkGroup",
        "athena:GetDataCatalog",
        "athena:GetDatabase",
        "athena:GetTableMetadata",
        "athena:ListDataCatalogs",
        "athena:ListDatabases",
        "athena:ListTableMetadata"
      ],
      "Resource": "*"
    }
  ]
}
```

```
}  
]  
}
```

Configure your Quick account to use the IAM role

After configuring the IAM role in the previous step, you must configure your Quick account to use it. For information about how to do that, see [Using existing IAM roles in Quick](#).

Update the identity propagation config with the AWS CLI

To authorize Quick to propagate end user identities to Athena workgroups, run the following `update-identity-propagation-config` API from the AWS CLI, replacing the following values:

- Replace `us-west-2` with the AWS Region that your IAM Identity Center instance is in.
- Replace `111122223333` with your AWS account ID.

```
aws quicksight update-identity-propagation-config \  
--service ATHENA \  
--region us-west-2 \  
--aws-account-id 111122223333
```

Create an Athena dataset in Quick

Now, create an Athena dataset in Quick configured with the IAM Identity Center enabled Athena workgroup you want to connect to. For information about how to create an Athena dataset, see [Creating a dataset using Amazon Athena data](#).

Key callouts, considerations, and limits

The following list contains some important considerations when using trusted identity propagation with Quick and Athena.

- Quick Athena data sources that use trusted identity propagation have Lake Formation permissions evaluated against the IAM Identity Center end user and the IAM Identity Center groups that the user might belong to.
- When using Athena data sources that use trusted identity propagation, we recommend any fine tuned access control is done in Lake Formation. However, If you elect to use Quick's scope down policy feature, scope down policies will be evaluated against the end user.

- The following features are disabled for data sources and data sets that use trusted identity propagation: SPICE datasets, Custom SQL on data sources, threshold alerts, email reports, Q Topics, stories, scenarios, CSV, Excel, and PDF exports, anomaly detection.
- If you experience high latency or timeouts, it may be because of a combination of high number of IAM Identity Center groups, Athena databases, tables, and Lake Formation rules. We recommend trying to use only the necessary number of those resources.

Data access integrations

Data access integrations in Amazon Quick establish secure connections to external data sources. They serve as the foundation for creating knowledge bases. Unlike connectors that perform actions, data access integrations focus on accessing and indexing content from third-party applications and services.

Data access integrations only configure authentication and point to the project or organization of the service. They cannot be used directly for analysis or by AI agents. You must create a connected knowledge base to make the data accessible.

How data access integrations work

Data access integrations configure authentication and establish connections to third-party service organizations or projects. You cannot use these integrations directly for analysis. You must create knowledge bases connected to the data access integration to make the data accessible for AI agents, chat interfaces, and spaces.

The relationship between data access integrations and knowledge bases is one-to-many:

- One data access integration can support multiple knowledge bases.
- Each knowledge base selects specific content from the connected data source.
- Knowledge bases inherit authentication and access permissions from their parent data access integration.

Create a data access integration

Use the following procedure to create a data access integration that establishes authentication and connection details for knowledge base creation. The following example demonstrates the process for setting up a Microsoft OneDrive data access integration, but the general steps are applicable to other data access integrations.

To create a data access integration

1. Scroll to the Set up a new integration area of the page. Find the application that you want to create an integration and knowledge base. Select "OneDrive".

Note

The Integration page defaults Knowledge bases tab and there may be existing knowledge bases that have been set up by others and shared. If you have previously set up an integration, check the data tab and use the action menu to create a knowledge base from there.

2. Select the plus (+) icon button on the application to create a new integration and knowledge base.
3. Select Bring data from Microsoft OneDrive option and click Next button.

Note

Some application integrations support data ingestion and read/write actions. The set up varies for each one. To set up actions, you'll need more information from your admin.

4. Complete the authentication process:
 - a. A Microsoft OneDrive sign in popup will appear automatically. If it doesn't, click Sign in to Microsoft OneDrive button.
 - b. Sign in using your Amazon credentials.
 - c. Wait until a success banner appears.
 - d. Click the "Next" button.
5. Select the data that should be ingested into the knowledge base using the file picker for OneDrive and click Add button.
6. Type in a knowledge base a Name and Description (optional), then click Create.
7. There will be a success toast notification and the data ingestion and sync will begin.
8. The data can take several minutes to sync, depending on the number files that are being ingested. The Status column will stay in the Syncing status until it is ready changing to Available.

9. When the knowledge base is ready, use the chat to ask questions and interact with it.

 Note

By default, the chat uses 'all data and apps' that you have access to and that are set up on your behalf. If you want to chat with a single knowledge base, select the knowledge base in the chat data picker.

 Note

You can also attach a knowledge base to a Space by navigating to the Space and adding it.

After successful creation, your data access integration appears in the integrations list. You can now create knowledge bases that use this integration to access and index content from the connected data source.

 Note

For detailed configuration steps specific to each data source, see [the section called "Supported integrations"](#).

Supported data sources

Amazon Quick supports data access integrations with the following applications and services. These integrations allow you to create knowledge bases from external data sources:

- **Amazon S3** - Access documents and files stored in S3 buckets using AWS credentials.
- **Atlassian Confluence** - Index pages, spaces, and attachments using user authentication or service authentication.
- **Google Drive** - Connect to personal and shared drives using OAuth 2.0 authentication.
- **Microsoft OneDrive** - Access OneDrive for Business content using user authentication or service authentication.

- **Microsoft SharePoint** - Index SharePoint Online and Server content using OAuth 2.0 authentication.
- **Web Crawler** - Index content from internal and external websites using basic authentication or form/SAML authentication.

Each data source supports different authentication methods and content access capabilities. The relationship between data access integrations and knowledge bases is one-to-many - one integration can support multiple knowledge bases, each selecting specific content from the connected data source.

Data source categories

Data access integrations are organized into the following categories based on the type of content and access patterns:

Cloud storage and file systems

- AWS S3 - Access documents and files stored in S3 buckets.
- Google Drive - Index content from personal and shared drives.
- Microsoft OneDrive - Connect to OneDrive for Business content.

Content management systems

- Atlassian Confluence - Access pages, spaces, and attachments.
- Microsoft SharePoint - Index SharePoint Online and Server content.

Web content

- Web Crawler - Index content from internal and external websites.

Authentication and security

Data access integrations use secure authentication methods to protect your data and maintain access controls. The authentication method depends on the specific data source and your organization's security requirements.

OAuth authentication

Most cloud-based integrations (Google Drive, OneDrive, Confluence Cloud) use OAuth 2.0 for secure, token-based authentication. This method allows Amazon Quick to access your data without storing your credentials.

Service account authentication

Enterprise integrations may use service accounts for programmatic access. This method is common for AWS S3 and other infrastructure-based data sources.

No authentication

Some integrations, such as web crawlers accessing public websites, may not require authentication. However, access controls are still enforced based on your Amazon Quick permissions.

Note

Authentication requirements and available methods vary by user tier. Readers may have limited authentication options compared to Authors.

Access control and permissions

Data access integrations maintain security by enforcing access controls at multiple levels. When users query content through knowledge bases, Amazon Quick ensures they can only access content they have permission to view.

- **Source-level permissions** - Users must have appropriate permissions in the source system (Google Drive, SharePoint, etc.).
- **Integration-level permissions** - Access to the integration itself is controlled by Amazon Quick permissions.
- **Knowledge base permissions** - Individual knowledge bases can have their own access controls.
- **Entity-level access controls** - When users query content, Amazon Quick verifies permissions for each document or item.

Key features and capabilities

Data access integrations provide several features to enhance your data integration experience:

- **Real-time synchronization** - Content is automatically updated when changes occur in the source system.

- **Selective indexing** - Choose specific folders, sites, or content types to include in your knowledge bases.
- **Content type support** - Index various file formats including documents, spreadsheets, presentations, and web pages.
- **Metadata preservation** - Maintain important metadata such as creation dates, authors, and tags.
- **Natural language querying** - Enable AI-powered search and question-answering across your indexed content.

Before you begin

Before creating data access integrations, ensure you have the following requirements in place:

- **Amazon Quick permissions** - Author or Admin role to create and manage integrations.
- **Source system access** - Appropriate permissions in the target system (administrative access may be required for some integrations).
- **Authentication credentials** - Valid credentials or service accounts for the target system.
- **Network connectivity** - Ensure Amazon Quick can access your data sources. Network requirements differ by integration type:
 - **Knowledge bases** - Do not support VPC connectivity. Data sources must be accessible over the public internet.
 - **Connectors** - Support VPC connectivity for resource servers within your VPC. However, authentication servers must remain publicly accessible.

Exploring your AWS data in Amazon Quick

Applies to: Enterprise Edition and Standard Edition

Intended audience: System administrators

Use this section to learn how to explore AWS data in Amazon Quick using the AWS Management Console. Using the **Explore in Amazon Quick Sight** shortcut, you can access a customizable dashboard template showing your data. Just as with any Amazon Quick Sight dashboard, this

dashboard can be refreshed on a schedule, published, and shared with other users in your organization.

To use this feature, you must first enable Amazon S3 analytics storage class analysis for your Amazon S3 buckets. For more on enabling storage class analysis in Amazon S3, see [Amazon S3 analytics – Storage class analysis](#) in the *Amazon Amazon S3 Developer Guide*.

After you have enabled storage class analysis, you can use Amazon Quick to explore your Amazon S3 analytics data.

To explore Amazon S3 analytics data in Amazon Quick

1. Open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. Choose a bucket to explore. The bucket must have storage class analysis enabled, with at least one filter.
3. Choose the **Management** tab.
4. Then choose **Analytics**.
5. Choose **Explore in QuickSight**.

Note

If you don't have an Amazon Quick account, you're prompted to create one before you can use the dashboard.

When you choose the option to explore in Amazon Quick, your Amazon S3 analytics data is automatically loaded into the dashboard template. The dashboard contains multiple visualizations to help you to understand the storage access pattern of your bucket.

Use the template as is, or customize it to suit your needs. For example, one visual on the default template helps you identify infrequently accessed data. It compares the amount of data retrieved to the amount of storage consumed, for objects of different ages.

You can also add your own visualizations to the dashboard. For example, you can break down the data access patterns, using filters for storage class analysis that you already have defined in Amazon S3 analytics.

To learn more about using S3 analytics and storage class analysis, see [Amazon S3 analytics – Storage class analysis](#) in the *Amazon Amazon S3 Developer Guide*.

AWS service connectors

With AWS service connectors in Amazon Quick, you can create connectors that interact directly with AWS services like Amazon Bedrock, Amazon Textract, and Amazon Comprehend. These connectors enable automated workflows that leverage AWS AI and machine learning capabilities.

What you can do

AWS service connectors enable you to integrate powerful AWS capabilities into your automated workflows. For example, you can use Amazon Bedrock to generate content with foundation models, Amazon Textract to extract text and data from documents, or Amazon Comprehend to analyze sentiment and extract insights from text. These connectors allow you to build sophisticated automation workflows that combine multiple AWS services for document processing, content generation, and data analysis—all while maintaining security through IAM role-based authentication.

Note

AWS services connectors can only be used with Amazon Quick Automate because they require an IAM identity for authentication. These connectors are created through the admin console and provide direct access to AWS service APIs.

Supported AWS services

Amazon Quick supports the following AWS services for connectors:

- **Amazon Bedrock Agent** - Invoke Bedrock agents for complex AI workflows.
- **Amazon Bedrock Runtime** - Access foundation models for text generation and conversation.
- **Amazon Bedrock Data Automation** - Automate data processing workflows with AI.
- **Amazon Comprehend** - Analyze text for sentiment, entities, and key phrases.
- **Amazon Comprehend Medical** - Extract medical information from healthcare text.
- **Amazon Textract** - Extract text and data from documents and images.
- **Amazon S3** - Manage objects and buckets in S3.

Before you begin

Before you set up an AWS service connectors, make sure you have the following:

- AWS account with access to the desired AWS services.
- IAM role with appropriate permissions for the AWS services you want to use.
- Amazon Quick Admin access to create connectors.
- Amazon Quick Automate access to use the connectors in workflows.

Prepare IAM role and permissions

Before setting up the connectors in Amazon Quick, prepare your IAM role with the necessary permissions for the AWS services you want to use.

Required IAM permissions

Configure your IAM role with the appropriate permissions based on the AWS services you plan to use. Make sure your IAM role includes a trust policy that allows Amazon Quick to assume the role for executing actions.

IAM Role for Resource Access

First, you will need to create an IAM role that will be used by Amazon Quick to call the AWS service needed in your Amazon Quick Automate workflow.

1. Log in to the AWS Console of the AWS account where the Amazon Quick subscription resides.
2. Open IAM and create a new IAM role.
3. Give it all the permissions for the AWS service you want to invoke via connectors. For example, you can assign a managed policy like `AmazonS3FullAccess` if you need to invoke Amazon S3.
4. In the trust relationship, give the assume role permission to `quicksight.amazonaws.com`. This allows Amazon Quick to assume this role and call AWS services on your behalf.
5. Once the Customer Role is created, take a note of the IAM role ARN.

Example trust policy:

```
{
  "Version": "2012-10-17"
  ,
  "Statement": [
    {
```

```
        "Sid": "Statement1",
        "Effect": "Allow",
        "Principal": {
            "Service": [
                "quicksight.amazonaws.com"
            ]
        },
        "Action": "sts:AssumeRole"
    }
}
```

Create AWS services connector

After preparing your IAM role and permissions, follow these steps to set up AWS Actions in Quick Suite.

Note

You need administrative access to Amazon Quick Suite to perform the following steps. See documentation on [Configure Amazon Quick subscriptions](#) to learn more.

Steps to create a new AWS integration

1. Choose the **profile icon**, then select **Manage Quick Suite**.
2. Select **AWS actions** under the **Permissions** section.
3. Choose **New Action** to create a new AWS integration.
4. Choose one of the **supported AWS services**.
5. Select **Next** to review available actions for this service.
6. Select **Next** to configure the connection details:
 - **Name** – Enter a descriptive name for your integration.
 - **Description** – Optionally, add notes about how this integration will be used.
 - **Role ARN** – Enter the ARN of the IAM role to be used for this AWS service.
7. Select **Next** to share the integration with users and/or user groups.
 - Provide **Owner access** for any users who need to edit, share, and delete the integration.

Note

Owner access is required to add integrations to Automation Groups in order to give access within Quick Automate.

- Provide **User access** for any users who need to invoke actions across Quick Suite.

For a list of integrations supported in various Quick Suite capabilities, see [Connector compatibility matrix](#).

8. Select **Add** to finish creating the integration.

See [Integration workflows](#) to learn more.

Next steps

After creating your action integration, you can:

- Share the integration with additional users or groups as needed.
- Add the integration to an **Automation Group** in order to use it in Quick Automate. See [Setup tasks](#) for more details.
- Monitor the integration's usage and performance through the admin console.
- Update the integration's configuration or permissions as requirements change.

Amazon Quick account management

If you are an Amazon Quick administrator, the account you use to sign in to Quick is in the ADMIN Quick group. There are also some permissions granted through IAM, which you might have already, or you can talk to your AWS account administrators to learn more.

Use the following topics to manage Quick.

Topics

- [Account details in Amazon Quick](#)
- [Amazon Quick asset management](#)
- [Manage Amazon Q permissions](#)

Account details in Amazon Quick

Use this section to change the account-wide settings in Amazon Quick. To open your account settings, choose the profile icon, and then select **Manage Quick**. Then, click on **Account settings**.

Deleting your Amazon Quick subscription and closing your account

Warning

Deleting your Amazon Quick subscription permanently destroys all organizational data across all AWS regions and cannot be reversed.

This action will immediately and permanently delete:

- Every dashboard, analysis, and dataset across your entire organization in all regions
- All user accounts and their associated work
- Custom data sources and connections
- All sharing permissions and folder structures
- Historical usage data and audit logs
- All customizations, themes, and organizational settings

Account deletion from any region affects your entire global Quick deployment. Even if you have Quick resources in multiple AWS regions, this single action deletes everything worldwide.

There is no recovery option. Once deleted, this data cannot be restored by AWS or recovered through any means.

Before proceeding: Export critical dashboards, back up datasets, and ensure all stakeholders are informed of this permanent action.

The act of deleting your Amazon Quick subscription is immediate and final. Deletion removes every Amazon Quick asset on the AWS account you are using. It doesn't delete namespaces that you added. (The **Default** namespace is deleted automatically.) You can locate and delete namespaces by using the API operations [ListNamespaces](#) and [DeleteNamespace](#).

You can terminate your Amazon Quick account from the **Manage Quick** menu or by using the API. To prevent someone from deleting a Amazon Quick user account accidentally or maliciously,

Amazon Quick uses permissions, a switch for the **Account termination protection** setting, and a required confirmation word.

After your account is deleted, you can create a new Amazon Quick account. The process doesn't take more than 15 minutes. The settings for edition and user authorization method on the new account can be the same or different.

Before you can delete your Amazon Quick account, make sure of the following:

- You're signed in using an IAM user or role with administrator permissions in Amazon Quick. We strongly recommend using an IAM identity instead of the AWS root account. For more information about root user security best practices, see [AWS account root user](#) in the *IAM User Guide*.
- You understand that your AWS account is not deleted when you terminate your Amazon Quick account. To instead close your AWS account, see [Closing an AWS account](#).
- Terminating your account deletes all users, all uploaded data, and assets (for example, datasets, data sources, queries, dashboards, analyses, settings, and so on).

Topics

- [Permissions and access to account termination](#)
- [Deleting your account](#)

Permissions and access to account termination

You need the following special permissions to terminate a Amazon Quick account. Without these permissions, you won't be able to terminate a Amazon Quick user account. Contact your account administrator for help.

- You're a Amazon Quick administrator and have an Admin role in Amazon Quick.
- You need permissions to run the following (except if you're the root admin user (IAM) who added Amazon Quick)
 - `quicksight:Unsubscribe`
 - `quicksight>DeleteAccountSubscription`
 - `ds:UnauthorizeApplication`
 - `ds>DeleteDirectory`

- `ds:DescribeDirectories`
- `quicksight:UpdateAccountSettings`

 Note

The `DeleteAccountSubscription` Amazon Quick API requires both `quicksight:Unsubscribe` and `quicksight>DeleteAccountSubscription` IAM permissions. An explicit deny on either permission results in an `AccessDeniedException`.

- To remove custom namespaces, you need permission to run the following API operations:
 - `quicksight:ListNamespaces`
 - `quicksight>DeleteNamespace`

You don't need extra permissions to delete the default namespace.

 Warning

Terminating your account is an instant action that cannot be undone by you or by AWS.

Deleting your account

The following sections outlines the different ways you can terminate your Amazon Quick instance.

To terminate your Amazon Quick account without the Amazon Quick UI

- Sign in to AWS where you want to remove Amazon Quick.

To terminate your account by using the Amazon Quick UI

1. Choose your profile on the application bar, and then choose **Manage Quick**.
2. Use one of the following methods to open the **Account termination** screen.
 - Use this [direct link](#) to the screen.
 - Choose **Account settings, Manage**.

3. On the **Account termination** page, confirm that you are viewing the correct Amazon Quick account by checking the name listed for account name.
4. Toggle off **Account termination protection is on**. Doing this enables the **Delete account** section.
5. For **Type "confirm" to delete this account**, enter the word confirmation word shown on your screen.

Amazon Quick asset management

Use this section to manage all of the assets in your Amazon Quick account in one unified view. You can manage the following Quick assets: Analyses, Dashboards, Datasets, Data sources, Shared folders, and Agents.

For the full model of how user removal affects a user's assets and data, see [User lifecycle and data handling in Amazon Quick](#).

Here are some common reasons to use the asset manager:

- **Transfer assets** – Quickly transfer assets from one user or group to another, for example when the original owner is no longer present.
- **Onboard new employees** – Speed up onboarding of new employees by giving them access to the same assets that their teammates using.
- **Support authors** – Better support authors in tenancies by giving support engineers temporary access to the author's dashboard.
- **Revoke access** – Quickly audit and revoke permissions, for example after implementations, customer support, or unexpected events.

To manage Amazon Quick assets

1. Choose the profile icon, then **Manage Quick**.
2. Open the asset manager by clicking on **Manage assets**.
3. You can search for assets by name, or browse for them in a list. Choose one of the following methods:

To search by name, select the appropriate search bar, using the name as your guide. Enter your search term and press **ENTER**. Find assets a user or group has access to by using the **Search by User or Group name** bar. Find other assets by using the **Search by asset name** bar.

To browse for assets by type, select a button by its name to view a type of asset, for example: browse analyses by selecting the **Analyses** button, browse for Data sources by selecting the **Data sources** button, and so on.

4. When you are viewing a list of your search results, you can interact with the assets listed. Here are a few examples:
 - Select an asset by toggling the box at the beginning of each row. Or, select everything by clicking the box at the top left of the list.
 - Change the type you are browsing for by selecting a different asset type from the **Asset type** list.
 - Use the vertical dot menu at right to perform an action on the asset in that row.
 - Use the **Share** button to share all of the selected assets. A popup window displays sharing options to apply to the users or groups you specify.
 - Select one or more assets, then choose the **Transfer** button to transfer the selected assets from one user or group to another. The **Transfer** button only appears after you select at least one asset. A popup window displays transferring options to apply to the users or groups you specify.

When you need to share a Amazon Quick asset with 100 or more users, consider using Amazon Quick groups. For more information on Amazon Quick groups, see [Creating and managing groups in Amazon Quick](#).

Manage Amazon Q permissions

Amazon Q is an AI-powered capability in Amazon Quick that enhances your data analysis experience. Quick admins can manage Amazon Q permissions in the Quick administration console. These permissions control how Amazon Q interacts with your data and users.

Topics

- [Manage personalization permissions](#)
- [Manage dashboard and visual indexing for search](#)
- [Manage Dashboard Q&A](#)

Manage personalization permissions

Amazon Q can use user metadata to provide more context-aware responses. This feature allows for a more personalized experience when interacting with Amazon Q.

To manage personalization permissions:

1. Open the [Quick console](#).
2. Choose the user icon at the top right, and then choose **Manage Quick**.
3. Under the **Account** section, choose **Amazon Q**.
4. Locate the **Manage personalization permissions** section.
5. Use the toggle switch next to **Personalized responses** to enable or disable this feature.

When enabled, Amazon Q can read user metadata to tailor its responses to individual users. This may include considering the user's role, location, or prior documents when providing information or suggestions.

Manage dashboard and visual indexing for search

Amazon Q can index information from dashboards to make them easily searchable across applications. This feature enhances the discoverability of your dashboards and their contents.

To manage dashboard and visual indexing:

1. Open the [Quick console](#).
2. Choose the user icon at the top right, and then choose **Manage Amazon Quick**.
3. Under the **Account** section, choose **Amazon Q**.
4. Locate the **Manage dashboard and visual indexing for search** section.
5. Use the toggle switch next to **Dashboard & visual indexing** to enable or disable this feature.

When enabled, Amazon Q can read dashboard metadata and visual information such as keywords, data fields, and refresh times to quickly respond to queries.

Manage Dashboard Q&A

The Dashboard Q&A feature allows Amazon Q to answer questions about your data without requiring the use of Quick Topics. This provides a more flexible and intuitive way for users to interact with their dashboards.

To manage Dashboard Q&A:

1. Open the [Quick console](#).
2. Choose the user icon at the top right, and then choose **Manage Amazon Quick**.
3. Under the **Account** section, choose **Amazon Q**.
4. Locate the **Manage Dashboard Q&A** section.
5. Use the toggle switch next to **Dashboard Q&A** to enable or disable this feature.

When enabled, Amazon Q can read dashboard metadata and visual information such as keywords, data fields, and refresh times to quickly respond to queries. This will eliminate the need for predefined Quick Topics.

Managing Quick subscriptions

Subscription management in Amazon Quick involves configuring and optimizing your service usage to meet your organization's needs while controlling costs. This section covers essential subscription-related tasks including upgrading or downgrading user subscriptions between Quick Professional and Enterprise, managing SPICE memory capacity to ensure optimal performance for your in-memory datasets, and configuring data capacity limits for unstructured data storage. Proper subscription management helps you balance performance, functionality, and cost while ensuring your users have the resources they need to create and share insights effectively.

Topics

- [Configure Amazon Quick subscriptions](#)
- [Configure SPICE memory capacity](#)
- [Index capacity](#)

Configure Amazon Quick subscriptions

You can purchase user subscriptions to get discounted pricing on Amazon Quick. When you invite additional users to Quick, you're charged for those accounts on a month-by-month basis.

To understand how Amazon Quick subscription names on the pricing page map to user roles in the admin console, see [Understanding Amazon Quick subscriptions and roles](#).

When you purchase an annual subscription, you pay for a Quick user account on an annual rather than monthly basis. With an annual subscription, you receive a discounted price in return for the

extended time commitment. You don't need to purchase an annual subscription to create or add users.

When you purchase a set of standard user subscriptions, you choose the number of accounts you want to cover. You also choose when to start the subscriptions (any time from the month following the current month, to one year in the future) and whether to autorenew them. All subscriptions that you purchase together must use the same values for these settings.

You can edit an existing set of user subscriptions to change whether it autorenews. If the set is not yet active, you can also change the number of subscriptions it covers, or delete it entirely.

Topics

- [Viewing current subscriptions](#)
- [Purchase subscriptions](#)
- [Editing a subscription](#)
- [Delete a subscription](#)
- [Upgrading or downgrading user subscriptions](#)
- [Understanding Amazon Quick user billing](#)

Viewing current subscriptions

Use the following procedure to view your current user subscriptions.

To view your current user subscriptions

1. Choose your user name on the application bar and then choose **Manage Quick**.
2. Choose **Manage pricing**.
3. Use the subscription meter to see how many accounts you have and how they are billed. In the following example, the account has 21 users total:
 - Seven users with annual subscriptions. Only currently active subscriptions are shown here.
 - 13 month-to-month users.

Pause over any section of the meter bar to display details about that user segment.

4. Use the information in the subscriptions table to see what current and future subscriptions you have.

Purchase subscriptions

Use the following procedure to purchase subscriptions.

To purchase subscriptions

1. Choose your user name on the application bar and then choose **Manage Quick**.
2. Choose **Manage pricing**.
3. Navigate to the **Authors and Admins** section, and then choose **Purchase plan**.
4. Choose or enter the number of subscriptions you want.
5. Choose the month and year when the subscriptions will start.
6. Choose whether the subscriptions autorenew.
7. Choose **Purchase**.

Editing a subscription

Use the following procedure to edit subscriptions.

To edit subscriptions

1. Choose your user name on the application bar and then choose **Manage Quick**.
2. Choose **Manage pricing**.
3. Next to the set of subscriptions you want to change, choose **Manage**.
4. (Optional) If the subscriptions haven't started yet, change the number of subscriptions.
5. Choose whether the subscriptions autorenew.
6. Choose **Save changes**.

Delete a subscription

Warning

Deleting Amazon Quick subscriptions affects user access. When you delete subscriptions:

- Users covered by deleted subscriptions may lose access to Quick features and content
- Subscription changes affect your entire organization's access to Quick capabilities

- Deleted subscriptions cannot be recovered - you must purchase new subscriptions to restore access

Before proceeding: Verify that affected users have alternative access methods and understand the implications of subscription changes.

Use the following procedure to delete subscriptions. You can only delete subscriptions that haven't started yet.

To delete subscriptions

1. Choose your user name on the application bar and then choose **Manage Quick**.
2. Choose **Manage pricing**.
3. Next to the set of subscriptions that you want to delete, choose **Edit**.
4. Choose **Delete Subscription**.

Note

If you use AWS Key Management Service or AWS Secrets Manager with Amazon Quick, you are billed for access and maintenance as described in the pricing pages for each AWS product. For more information on how these products are billed, see the following:

- [AWS Key Management Service Pricing page](#)
- [AWS Secrets Manager Pricing page](#)

In your billing statement, the costs are itemized under the appropriate product and not under Amazon Quick.

Upgrading or downgrading user subscriptions

Amazon Quick Professional and Amazon Quick Enterprise are per-user subscriptions that you assign through group membership. To change a user's subscription, you move them between the groups that are mapped to Quick roles. The Reader Pro role corresponds to the Amazon Quick Professional subscription, and the Author Pro role corresponds to the Amazon Quick Enterprise

subscription. For more information about how subscriptions map to roles, see [Understanding Amazon Quick subscriptions and roles](#).

For a full comparison of the capabilities included with each subscription and current per-user pricing, see [Amazon Quick pricing](#).

Upgrading from Professional to Enterprise

To upgrade a user from the Amazon Quick Professional subscription to the Amazon Quick Enterprise subscription, move the user to a group that is mapped to the Author Pro role. If you use IAM Identity Center or Active Directory, you make this change by updating the user's group membership in your identity provider. For the procedure, see [Changing a user's role](#).

Users can also request an upgrade themselves when they encounter a feature or usage limit that requires the Enterprise subscription. Depending on your organization's settings, these requests can process automatically or require administrator approval. For more information, see [User-driven license upgrades](#).

After the change takes effect, the user has immediate access to Enterprise capabilities. Changes to users or groups can take up to five minutes to propagate.

Downgrading from Enterprise to Professional

To downgrade a user from the Amazon Quick Enterprise subscription to the Amazon Quick Professional subscription, move the user out of the group that is mapped to the Author Pro role and into a group that is mapped to the Reader Pro role. As with an upgrade, you make this change through your identity provider when you use IAM Identity Center or Active Directory. For the procedure, see [Changing a user's role](#).

Important

If a user belongs to multiple groups that are mapped to different roles, they keep the subscription that grants the broadest level of access. To complete a downgrade, remove the user from every group that is mapped to the Author Pro role.

After you downgrade a user, they lose access to the capabilities included only with the Amazon Quick Enterprise subscription. Before you downgrade, make sure any assets the user owns that you want to keep are transferred to another user. Changes to users or groups can take up to five minutes to propagate.

Note

For questions about billing adjustments or refunds related to a subscription change, contact [AWS Support](#).

Understanding Amazon Quick user billing

Your charges for a Quick user follow the user's lifecycle in your account, from the moment you provision the user through the period after you remove the user. The following sections explain when billing begins, how your first billing period is prorated, how ongoing monthly charges apply, and how removing a user affects your charges.

Topics

- [When billing begins](#)
- [Your first billing period](#)
- [Ongoing monthly billing](#)
- [Billing when you remove a user](#)
- [Billing examples](#)

When billing begins

Billing for a Quick user begins when you provision the user in your account, not when the user first signs in. This is independent of user activity. If you provision a user who never signs in during a billing period, you are still billed for that user for that period.

The exact moment billing begins depends on how you provision the user:

- For users you provision through the API, such as by registering a user or creating an account subscription, billing begins when the API call creates the user, whether or not the user has signed in.
- For users who self-provision, billing begins the first time the user opens Amazon Quick.
- For users you map from a group in IAM Identity Center or Active Directory, billing begins when Amazon Quick detects the user through group synchronization.

Your first billing period

When you provision a user partway through a billing period, your charge for that first period is prorated. The proration is based on the portion of the period that remains after the provisioning date. In the next full billing period, you pay the standard rate for the user's role.

Ongoing monthly billing

For each billing period a user remains in your account, you pay the monthly rate for that user's role. These subscription charges are separate from usage-based charges such as SPICE capacity or reader session usage. For a full comparison of the capabilities included with each subscription and current per-user pricing, see [Amazon Quick pricing](#).

Billing when you remove a user

When you remove a user partway through a billing period, you still pay for the full billing period in which the removal occurs. You are not charged for the user in any later billing period. This differs from provisioning. Your first billing period is prorated, but your final billing period is not.

To avoid charges for a user you no longer need, remove the user before the next billing period begins.

Billing examples

The following examples show how your charges apply across a user's lifecycle:

- For a user you provision partway through a billing period, your charge for that period covers the provisioning date through the end of the period, and the full rate applies in the next period.
- For a user you remove partway through a billing period, you pay for the full period, and no charge applies in the following period.
- For a user you provision and remove within the same billing period, you pay for that full period, and no charge applies afterward.
- For a user who never signs in, you pay each billing period the user remains in your account, the same as an active user, until you remove the user.

This billing behavior applies to all Amazon Quick user roles across editions and in all AWS Regions where Amazon Quick is available.

Configure SPICE memory capacity

SPICE (Super-fast, Parallel, In-memory Calculation Engine) is the robust in-memory engine that Amazon Quick uses. It's engineered to rapidly perform advanced calculations and serve data. In Enterprise edition, data stored in SPICE is encrypted at rest. For more information, see [Data encryption in Quick](#).

SPICE capacity is allocated separately per AWS Region. For each AWS account, SPICE capacity is shared by all the people using Quick in a single AWS Region. The other AWS Regions have no SPICE capacity unless you choose to purchase some.

Quick administrators can view how much [SPICE](#) capacity you have in each AWS Region and how much of it is currently in use. Administrators can also purchase additional SPICE capacity or release unused SPICE capacity. You can only release SPICE capacity that isn't currently used by a dataset. Datasets in SPICE stay there until someone remove them from SPICE. To change that, you can either delete the datasets or change them so they aren't stored in SPICE.

Purchasing or releasing SPICE capacity only affects the capacity for the currently selected AWS Region. Each AWS account can have a separate Amazon Quick subscription and can be used in multiple AWS Regions.

Before you make any changes to SPICE capacity, make sure that you're using the correct AWS account and AWS Region. It's possible to be using different AWS accounts or AWS Regions at the same time in different contexts, as follows:

- If you open Amazon Quick using the `http://quicksight.aws.amazon.com` URL, Amazon Quick automatically selects your account and AWS Region. You can't view your AWS account from Amazon Quick. We recommend using a different method to open Amazon Quick when you want to work with SPICE capacity.
- If you open Amazon Quick from the AWS Management Console, Amazon Quick opens in the account that you used to sign in to that console. However, it opens in the last AWS Region that you selected in Amazon Quick. The AWS Management Console and the Amazon Quick console each have an AWS Region selector that works independently from the other. Changing the selected AWS Region in the AWS console doesn't change the AWS Region in Amazon Quick.
- If you use the AWS Command Line Interface (AWS CLI) to run Amazon Quick commands, make sure to provide the relevant AWS account for each Amazon Quick API operation you use. The AWS Region isn't always required, and if you don't provide it, the AWS CLI uses your default AWS Region from your AWS configuration. We recommend that you always explicitly provide the AWS Region, to make sure you apply the command to the correct AWS Region.

You must be signed in as a Amazon Quick administrator to view or manage SPICE capacity.

Topics

- [Finding your current AWS account and AWS Region](#)
- [Viewing SPICE capacity and usage in an AWS Region](#)
- [Hiding SPICE capacity labels](#)
- [Purchasing SPICE capacity in an AWS Region](#)
- [Turning on SPICE auto capacity purchasing](#)
- [Releasing SPICE capacity in an AWS Region](#)

Finding your current AWS account and AWS Region

To select the correct AWS account and AWS Region (console)

1. Open the AWS console, using the AWS account that you want to view SPICE information for. If you have only one AWS account, you can skip this step.

You can verify the account number by following these steps:

- a. On the navigation bar at the top of the page, choose the account name or number at right. If a number displays, this might be your AWS account ID.
- b. Choose **My Security Credentials** to display your credential-related information and options. Your AWS account ID displays near the top of the page.

To return to the original page, choose the AWS logo at upper left.

2. Open Amazon Quick by first entering "**quicksight**" into the **Find Services** search box. When the word Amazon Quick appears following the search box, choose it from the list.
3. In Amazon Quick, open the profile menu by choosing your profile icon at top right. The AWS name of the AWS Region that Amazon Quick is using displays in the menu.

The same AWS Region also displays in the URL, for example: `https://us-east-1.quicksight.aws.amazon.com/sn/admin`. If this is your URL, the profile menu displays the name N. Virginia.

To switch AWS Regions, display the list of supported Regions by choosing the Region name from the profile menu. Then choose the Region that you want to use. Switching to a different

AWS Region changes the SPICE usage information that you can view. It also changes the Amazon Quick assets that you can use, for example data sources and dashboards.

Viewing SPICE capacity and usage in an AWS Region

To view current SPICE capacity and usage (console)

1. Open Amazon Quick. Make sure that you're using the correct AWS account and AWS Region as described previously in [Finding your current AWS account and AWS Region](#).
2. Open the administration page by choosing **Manage Quick** from your profile menu.
3. Choose **SPICE capacity** from the left navigation pane . The following information displays:
 - The **Total SPICE capacity** section displays the total amount of used and unused SPICE capacity. A bar graph shows how much of this storage space is in each of the following categories for this AWS account in the AWS Region that's currently selected in Amazon Quick:
 - Purchased SPICE capacity – This is the additional SPICE capacity.
 - SPICE capacity bundled with Amazon Quick – This is the total default capacity associated with your paid users.

Hover over any section of the meter to see details on that capacity type.

- The **SPICE usage** section displays the total amount of the used and unused SPICE capacity. A bar graph shows how much of this storage space is in each of the following categories for this AWS account in the AWS Region that's currently selected in Amazon Quick:
 - Used SPICE capacity – This is the used portion of the default SPICE capacity that you get for each user.
 - Unused SPICE capacity – This is the unused portion of the default SPICE capacity that you get for each user.
 - Releasable unused capacity – This is the purchased capacity that isn't in use, and so can be released to reduce costs.

Hiding SPICE capacity labels

Amazon Quick account admins can choose to hide the account-wide SPICE capacity usage and remaining size labels from Amazon Quick authors. This feature is available to all enterprise

accounts that use custom permissions. For more information about custom permissions in Amazon Quick, see [Customizing access to Amazon Quick capabilities](#)

Use the following procedure to hide SPICE capacity usage from the Amazon Quick console.

1. Open the [Quick console](#).
2. From any page in the Amazon Quick console, choose your profile name, and then choose **Manage Quick**.

The **Manage Quick** menu is only available to Amazon Quick admins. If you are unable to access this menu, contact your Amazon Quick account admin for assistance.

3. Choose **Manage users**, and then choose **Manage permissions**.
4. Edit or create a new custom permission. For **Restrict access to**, under **Datasets**, select **Viewing account SPICE capacity**.
5. When you are finished creating or changing the custom permission, choose **Create** or **Update**.

After you create or update a custom permission to hide SPICE capacity usage, assign the new permission to users with the `UpdateUser` API.

Purchasing SPICE capacity in an AWS Region

To purchase more SPICE capacity (console)

1. Open Amazon Quick. Make sure that you're using the correct AWS account and AWS Region as described previously in [Finding your current AWS account and AWS Region](#).
2. Open the administration page by choosing **Manage Quick** from your profile menu.
3. Choose **SPICE capacity** from the left navigation pane .
4. Choose the **Purchase more capacity** button.
5. Enter a number of gigabytes of SPICE capacity to purchase for the AWS Region that is currently selected in Amazon Quick.
6. To confirm your choice, choose **Purchase SPICE capacity**. To exit without making any changes, choose **Cancel**.

Turning on SPICE auto capacity purchasing

Turn on SPICE auto capacity purchasing to allow Amazon Quick to automatically manage your Amazon Quick account's SPICE capacity. When you turn auto capacity purchasing on, Amazon

Quick evaluates how much capacity is needed based on your account's usage. As your account uses more SPICE storage, Amazon Quick automatically purchases SPICE capacity as needed on your behalf. This allows users to ingest data as needed without the need to estimate usage or manually purchase SPICE data. Auto capacity purchasing makes it easier for new customers, ISVs, and larger companies to directly access SPICE without needing to understand, track, or manually purchase their account's SPICE capacity. Amazon Quick admins can still purchase and release SPICE capacity manually.

Auto capacity purchasing doesn't support auto-decrement. If users want to reduce their SPICE usage, capacity release must be done manually.

By default, all new Amazon Quick accounts that are created in the Amazon Quick console have auto capacity purchasing turned on in the region that their capacity is located. To turn on auto capacity purchasing for other regions, Amazon Quick account admins can manually turn on auto capacity from the **SPICE capacity** management page.

By default, all new Amazon Quick accounts that were created with the Amazon Quick API and all existing Amazon Quick accounts have auto capacity purchasing turned off. To turn on auto capacity purchasing, Amazon Quick account admins can manually turn on auto capacity from the **SPICE capacity** management page.

To turn SPICE capacity purchasing on or off

1. Open the [Quick console](#).
2. From any page in the Amazon Quick console, choose your profile name, and then choose **Manage Quick**.

The **Manage Quick** menu is only available to Amazon Quick admins. If you are unable to access this menu, contact your Amazon Quick account admin for assistance.

3. Choose **SPICE capacity**.
4. On the **SPICE Capacity** page that opens, toggle the **Auto-purchase capacity** on.

To turn auto capacity purchasing off, follow the procedure above and toggle **Auto-purchase capacity** off. When auto purchase capacity is turned off, ingestions or refreshes that exceed the account's SPICE capacity automatically fail.

Amazon Quick admins can turn auto capacity pricing on or off at any time. If you turn auto capacity purchasing off after it's been in use, your account's current capacity becomes your account's

purchased capacity. If your account has no remaining capacity when you turn auto purchase off, the next ingestion or refresh will fail.

If your account already exceeds its SPICE capacity when you turn auto capacity purchasing on, Amazon Quick automatically matches your account's capacity to your current usage. After Amazon Quick matches your account's capacity, the auto-purchase logic starts.

Releasing SPICE capacity in an AWS Region

To release unused SPICE capacity (console)

1. Open Amazon Quick. Make sure that you're using the correct AWS account and AWS Region as described previously in [Finding your current AWS account and AWS Region](#).
2. Open the administration page by choosing **Manage Quick** from your profile menu.
3. Choose **SPICE capacity** from the left navigation pane .
4. Choose **Release unused purchased capacity**.
5. Do one of the following:
 - To release all SPICE capacity from the AWS Region that is currently selected in Amazon Quick, choose **Release all**.
 - To release some of gigabytes of SPICE capacity from the AWS Region that is currently selected in Amazon Quick, enter the number of gigabytes to release.
6. To confirm your choice, choose **Release SPICE capacity**. To exit without making any changes, choose **Cancel**.

Index capacity

An index is a structured corpus of documents that allows for efficient searching, sorting, and data access. Amazon Quick uses an index to store, prepare, sync any files you upload to your Quick spaces. It then uses this indexed data to generate agent-driven responses during chat conversations.

Topics

- [Storage allocation](#)
- [Regional allocation](#)
- [Manual scaling & autoscaling](#)

- [Total index capacity](#)
- [To view index capacity usage](#)
- [Release unused index capacity](#)
- [Switch between capacity modes](#)

Storage allocation

You get index capacity measured as the original file size of documents at their source location (such as SharePoint, S3, Confluence) before any processing by Quick. Quick automatically creates an index when you sign up in your home region with auto-scaling enabled (a mechanism where your index capacity is adjusted automatically based on usage). Your Quick Index storage allocation is based on your user count and subscription tier. **Professional users** get 25 GB of storage per user and **Enterprise users** get 50 GB per user. Overages are billed at \$5 per GB per month. Allocated storage and overages are pooled on a payer account level. This replaces the previously auto-provisioned 50 MB (extracted text) default index capability.

Regional allocation

Your Quick Index storage allocation is billed as a discount to your home region, the region you selected during your Quick subscription setup. If you provision Quick Index capacity in additional regions beyond your home region, that storage counts as overage and is billed at \$5 per GB per month.

Manual scaling & autoscaling

You can purchase additional capacity via the Quick console in two ways. First, auto-purchase is enabled by default for all accounts. When your usage approaches your organizational allocation, the system automatically purchases additional capacity in 0.25 GB increments up to a maximum of 2 TB. Second, you can manually purchase capacity from 0.25 GB, up to 2 TB in 0.25 GB increments. The pre-populated auto-purchase limit is set to 2 TB, which you can adjust in the console. The default manual purchase amount is set to 0.25 GB. For manual purchase mode, you will be billed based on the storage capacity you have purchased, while in auto scale mode you will be billed based on 0.25 GB increment bundles, depending on how much you have consumed.

Quick Index has an auto-scaling limit currently set at 2 TB of storage per account, replacing the previous limit of 60 GB of extracted text. If you need to index more than 2 TB, please open a ticket with AWS support to extend this limit.

Total index capacity

You must purchase data storage capacity for your Amazon Quick index to ingest data from knowledge bases and upload files to spaces. In the index capacity page, the first horizontal bar shows your total index capacity. This is the current capacity available in your index to support file uploads and knowledge bases. You can switch between manual and auto-scaled index capacity modes as per your needs. When auto-scaling mode is **ON**, this capacity automatically adjusts to reflect the latest available capacity based on your ingestion needs. When manual mode is **ON**, this total capacity represents the actual capacity you have purchased for your account based on your estimate of the ingestion needs.

If you previously purchased Index capacity measured in extracted MBs of storage, your existing purchased capacity will be converted to the new Raw Storage dimension automatically. When buying capacity manually, you can also release unused capacity. If you want to release all provisioned index capacity, you can do so after deleting corresponding data—like file uploads in spaces and knowledge bases—in your Amazon Quick instance.

Note

With auto-scaling enabled, Amazon Quick may temporarily provision capacity above your maximum purchase limit during periods of high concurrent usage to maintain performance. However, you are only charged for capacity up to your configured maximum purchase limit, not for any temporary over-provisioning that occurs.

To manage purchase options

1. Log in to the Amazon Quick console.
2. Select **Manage Quick**.
3. From the admin console left navigation menu, select **Subscriptions**, and then select **Index capacity**.
4. In **Index capacity**, from **Total index capacity** select **Manage capacity**, and do the following:
 - a. In **Manage purchase options**, do one of the following:
 - i. To automatically scale index capacity as per usage, select **Auto-scale (Recommended)**. This automatically adjusts index capacity in 0.25 GB increments based on your usage, purchasing or releasing as needed. Then, in **Maximum purchase**

limit set the limit at which auto-purchase stops. The value of the maximum purchase limit must be 2 TB. Maximum purchase limit cannot be lower than current used capacity.

- ii. To manually purchase capacity, select **Manual**. Then, in **Index capacity**, set the total capacity you want your index to hold. The capacity must be between 0.25 GB-2 TB and you can increase it in 0.25 GB increments. When you want to lower the current capacity held by the index, it can't be lowered below the used capacity. Since capacity is managed in 0.25 GB increments, you can only decrease to the next 0.25 GB unit above your current usage. For example, if your index has 2 GB capacity with 1.1 GB consumed, you can reduce capacity to 1.25 GB but not to 1 GB since part of the 1-1.25 GB unit is already in use.
- b. Select **Confirm**.

To view index capacity usage

The Amazon Quick admin console displays both your purchased index capacity and your index capacity usage. The following procedure shows you how to view them.

To view index data capacity usage

1. Log in to the Amazon Quick console.
2. Select **Manage Quick**.
3. From the left navigation menu, select **Subscriptions**, and then select **Index capacity**.
4. In **Index capacity**, you'll find your capacity usage under **Index capacity usage**. First bar shows you total purchased capacity and second bar shows you index capacity usage.

Release unused index capacity

The Amazon Quick admin console displays both your purchased index capacity and your index capacity usage. If you've manually provisioned capacity, you can choose to release unused index capacity. If you've enabled auto-scaling for your index, you don't need to release index capacity as Amazon Quick automatically scales index capacity to your usage. The following procedure shows you how to release unused index capacity.

 Note

To release any provisioned index capacity, you must delete all datasets and files uploaded to spaces. Then, follow the steps to release unused index capacity.

To release index data capacity

1. Log in to the Amazon Quick console.
2. Select **Manage Quick**.
3. From the left navigation menu, select **Subscriptions**, and then select **Index capacity**.
4. In **Index capacity**, select **Manage capacity**.
5. In **Manage index capacity**, reduce the **Index capacity** to your desired value.

Switch between capacity modes

You can switch between manual and auto-scale index capacity modes at any time. When switching from auto-scale to manual, your current capacity becomes your fixed manual capacity. When switching from manual to auto-scale, you set a maximum purchase limit and Amazon Quick automatically adjusts capacity based on usage.

To switch from auto-scale to manual capacity

1. Log in to the Amazon Quick console.
2. Select **Manage Quick**.
3. From the admin console left navigation menu, select **Subscriptions**, and then select **Index capacity**.
4. In **Index capacity**, from **Total index capacity** select **Manage capacity**.
5. In **Manage index capacity**, select **Manual**.
6. In **Index capacity**, set your desired manual capacity. The value must be between 0.25 GB-2 TB in 0.25 GB increments and can't be less than current usage.
7. Select **Confirm**.

To switch from manual to auto-scale capacity

1. Log in to the Amazon Quick console.

2. Select **Manage Quick**.
3. From the admin console left navigation menu, select **Subscriptions**, and then select **Index capacity**.
4. In **Index capacity**, from **Total index capacity** select **Manage capacity**.
5. In **Manage index capacity**, select **Auto-scale (Recommended)**.
6. In **Maximum purchase limit**, set your maximum purchase limit. The value must be between 0.25 GB-2 TB in 0.25 GB increments and can't be less than current usage.
7. Select **Confirm**.

Managing identities

Identity management in Amazon Quick involves controlling who can access your Quick resources and what actions they can perform. This section covers the essential tasks for managing user identities, including adding and removing users, organizing them into groups for easier permission management, and transitioning between different identity management systems. Whether you're using IAM Identity Center, federated identity providers, or other authentication methods, these topics will help you establish proper access controls and maintain security for your Quick environment.

Before managing user access, it's important to understand how Amazon Quick subscription names map to user roles. For information about subscription types and their capabilities, see [Understanding Amazon Quick subscriptions and roles](#).

Topics

- [Managing user access inside Amazon Quick](#)
- [User-driven license upgrades](#)
- [Creating and managing groups in Amazon Quick](#)

Managing user access inside Amazon Quick

Intended audience: System administrators and Amazon Quick administrators

Amazon Quick administrators manage user access in Quick. User access management in Quick is determined by your Quick account identity configuration. For accounts that use IAM Identity Center

or Active Directory, groups are assigned to Quick roles. Groups can be assigned the Admin, Author, Reader, Admin Pro, Author Pro, or Reader Pro roles. For more information about Pro roles in Quick see [Get started with Generative BI](#). Note that Reader Pro maps to Amazon Quick Professional subscription and Author Pro maps to Amazon Quick Enterprise subscription. For more information about integrating your Amazon Quick account with IAM Identity Center, see [Managing access for IAM Identity Center users](#).

To understand how subscription names relate to user roles and their capabilities, see [Understanding Amazon Quick subscriptions and roles](#).

Amazon Quick accounts that use Quick and IAM users create users directly in Quick. These users and their roles are managed at the user level. For more details, see [Managing access for Quick and IAM users](#).

Topics

- [Managing access for IAM Identity Center users](#)
- [Managing access for Amazon Quick and IAM users](#)

Managing access for IAM Identity Center users

Applies to: Enterprise Edition

Intended audience: System administrators and Amazon Quick administrators

AWS administrators can use this topic to learn more about managing accounts that are integrated with IAM Identity Center. The information in this section also applies to Quick accounts that use Active Directory.

To manage Quick users, you must have administrative privileges in Quick and also the appropriate AWS permissions. For more information about the necessary AWS permissions, see [IAM policy examples for Quick](#). If you are using directory groups, you need to be a network administrator.

Each Quick Enterprise edition account can have an unlimited number of users. User names that contain a semicolon (;) aren't supported.

Use the following topics to add, view, and deactivate Quick users.

Important

You can't remap Amazon Quick users or groups from one identity store to another. For example, if you are migrating from an on-premises Active Directory to Directory Service, or the other way around, you unsubscribe and resubscribe to Amazon Quick. You do this because even if the user's aliases remain the same, the underlying identity data changes. To make the transition easier, request in advance that your users document all their Amazon Quick assets and settings before the migration.

Topics

- [Adding users](#)
- [Managing user access](#)
- [Deactivating user accounts](#)
- [Changing a user's role](#)
- [Deleting Enterprise accounts](#)

Adding users

With IAM Identity Center, add users to Amazon Quick by associating their IAM Identity Center group to an Admin, Admin Pro, Author, Author Pro, Reader, or Reader Pro role in Quick. All users in the selected groups are authorized to sign in to Quick.

For more information about Pro roles in Quick see [Get started with Generative BI](#).

To see which groups are integrated with your Quick account, follow the procedure in [Managing user access](#).

Managing user access

Use the following procedure to view groups that are assigned to a role that grants access to Quick.

1. Open the [Quick console](#).
2. Choose your user name on the application bar and then choose **Manage Account**.
3. Choose **Manage Users**.
4. Choose **Manage role groups**.

5. In the **Manage role groups** page, use the tables to add or remove groups in IAM Identity Center or Active Directory from the Admin, User, or Reader roles in Quick.

Deactivating user accounts

Deactivating a Amazon Quick group or user account removes that group or user's access to Quick resources, like analyses or data sets. IAM Identity Center or Active Directory users that are removed from a group that grants them access to Quick lose access to Quick. These users appear in the **Inactive users** list in Quick and remain there until an administrator removes them. Before you deactivate a user, you can reassign their resources to another user with the asset management console.

If you later need to reactivate a Quick user's account, put the user into a group with access to Quick. Doing this restores their access to Quick and to any existing resources that are still associated with that user.

Note

With IAM Identity Center integrated into your Amazon Quick account or Active Directory users, you can change a user's role type by moving them to a group that is associated with a different Amazon Quick role. If a user is in multiple groups that are mapped to different Amazon Quick role types, the user is able to access Amazon Quick with the role that offers the broadest level of access. Accounts that use other identity types can't upgrade or downgrade a user by transferring them between groups. For more information, see [Changing a user's role](#).

You can activate or deactivate multiple users at once by adding or removing one or more IAM Identity Center or Active Directory groups that are associated with a role in Amazon Quick.

Changing a user's role

If you're using IAM Identity Center or Active Directory, you can change a user's role by adding or removing them from a group that's mapped to the role that you want to assign them in Quick. You can also perform this task by adding a new group to a role in Quick. To do this, you need both administrative privileges in Quick and also appropriate AWS permissions.

With IAM Identity Center integrated users, you can change role types for a user by moving them to a group that is associated with a different Quick role. If a user belongs to multiple groups that

are mapped to different role types, the user is able to access Quick with the role that offers the broadest level of access.

When you make changes to users or groups in Quick, it can take up to five minutes for the change to take effect. Examples of such changes are the following:

- Deleting a user
- Changing a user from an admin to an author
- Adding or removing group members

The five-minute time period allows changes to propagate throughout the system.

Deleting Enterprise accounts

If a user is deleted from IAM Identity Center or Active Directory or is removed from a group that's associated with a role in Quick, the user no longer exists in Quick. You do not need to delete the user in the Quick application. The deleted user will appear in the **Inactive users** list in Quick until the first day of the following month. After that date passes, the user is automatically removed from the list.

When a user is removed from your identity provider, the user's assets are not transferred or deleted automatically. Administrators can review deleted users and transfer their assets to another user. To transfer assets that no longer have an owner, see [Amazon Quick asset management](#). For the full model of how user removal affects assets and data, see [User lifecycle and data handling in Amazon Quick](#).

Removing a user affects billing for the current billing period. For how billing works when you remove a user, see [Understanding Amazon Quick user billing](#).

Managing access for Amazon Quick and IAM users

Amazon Quick account administrators can use this topic to learn more about managing accounts that use IAM or Quick for identity federation.

To manage Quick users, you must have administrative privileges in Quick and also the appropriate AWS permissions. For more information about the necessary AWS permissions, see [IAM policy examples for Quick](#).

Each Quick Enterprise edition account can have an unlimited number of users. User names that contain a semicolon (;) aren't supported.

Use the topics below to learn more about managing access for Quick and IAM users.

Topics

- [Inviting users to access Amazon Quick](#)
- [Viewing Amazon Quick account details](#)
- [Deleting a Amazon Quick user account](#)

Inviting users to access Amazon Quick

Applies to: Enterprise Edition and Standard Edition

Intended audience: Amazon Quick administrators

Use the following procedure to invite a user to access Quick.

1. Choose your user name on the application bar and then choose **Manage Quick**.
2. Choose **Manage Users**. On this screen, you can manage users who already exist in your account.
3. Choose **Invite users**.
4. In the **Invite users to this account** table, enter a new user name for a person to whom you want to grant access to Quick. If the user is an IAM user, enter their IAM username. Then press **+**. A user's IAM username can be the same as their email address.

Repeat this step until you have entered everyone who you want to invite. Then go to the next step to enter details.

5. For **Email**, enter an email address for the account.

Note

Currently, email addresses are case-sensitive. Email distribution lists are not supported for inviting users. Instead, send individual invitations to each user. To share assets like dashboards, datasets, and agents with multiple users, create user groups in Quick.

6. For **Role**, choose the role to assign to each person you're inviting. A *role* determines the permission level to grant to that account.

- **ADMIN roles:**

- **ADMIN** – The user is able to both use Amazon Quick for authoring and for performing administrative tasks like managing users or purchasing [SPICE](#) capacity.
- **ADMIN PRO** – The user is able to perform all actions of a Amazon Quick Admin and utilize applicable Amazon Quick Generative BI capabilities. For more information about Pro roles in Amazon Quick, see [Get started with Generative BI](#).

There are some differences in the administrative tasks that IAM users and Amazon Quick administrators can perform. These differences occur because some administrative tasks require permissions in AWS, which Amazon Quick-only users lack. The differences are these:

- Amazon Quick administrators can manage users, SPICE capacity, and subscriptions.
- IAM users with administrative permissions can also manage users, SPICE capacity, and subscriptions. In addition, they can manage Amazon Quick permissions to AWS resources, upgrade to Enterprise edition, and unsubscribe from Amazon Quick.

If you want to create a user with administrator permissions with IAM access, check with your AWS administrator. Make sure that the IAM user has the all necessary statements in their IAM permissions policy to work with Amazon Quick resources. For more information about what statements are required, see [IAM policy examples for Amazon Quick](#).

- **AUTHOR roles:**

- **AUTHOR**– The user is able to author analyses and dashboards in Amazon Quick but not perform any administrative tasks in Amazon Quick.
- **AUTHOR PRO**– The user is able to perform all actions of a Amazon Quick Author and utilize applicable Amazon Quick Generative BI capabilities. For more information about Pro roles in Amazon Quick, see [Get started with Generative BI](#).

- **READER roles (Enterprise only):**

- **READER**– Users are able to interact with shared dashboards, but not author analyses or dashboards or perform any administrative tasks.
- **READER PRO**– The user is able to perform all actions of a Amazon Quick Reader and utilize applicable Amazon Quick Generative BI capabilities. For more information about Pro roles in Amazon Quick, see [Get started with Generative BI](#).

7. For **IAM User**, verify that it says **Yes** for accounts that are associated with IAM users, and **No** for those that are Amazon Quick-only.
8. (Optional) To delete a user, choose the delete icon at the end of the relevant row.
9. Choose **Invite**.

Viewing Amazon Quick account details

Intended audience: Amazon Quick administrators

You can view Amazon Quick accounts on the **Manage Users** page. To view a Amazon Quick user account, use the following procedure.

1. Choose your user name on the application bar and then choose **Manage Quick**.
2. Choose **Manage Users** to view details about people who are Amazon Quick users. The information that displays includes:
 - Username – The person's user name.
 - Email – The email associated with this user name.
 - Role – The security cohort that the person's user name belongs to: **ADMIN**, **ADMIN PRO**, **AUTHOR**, **AUTHOR PRO**, **READER**, or **READER PRO**.
 - Last active – The last date and time that this person accessed the Amazon Quick console. Anyone who isn't an active user has a **Last active** status of User has no activity.

You can also see deleted or inactive users in this screen.

3. To find a user name, enter a part or all of a user's name or email the search box. Search is case-insensitive and wildcards aren't supported. To clear the search results and view all user names, delete your search entry.

Deleting a Amazon Quick user account

Intended audience: Amazon Quick administrators

⚠ Warning

Deleting user accounts has permanent, organization-wide consequences. When you delete a user account:

- Assets owned solely by the user are permanently removed unless you transfer them to another user before deletion.
- Assets that have other owners remain available to those owners; the deleted user is removed from their permissions.
- **This action cannot be undone – deleted resources cannot be recovered.**

Always transfer critical resources to another user before deleting an account. Review all user-owned assets using the asset management console before proceeding with any account deletion.

For the full model of how user removal affects assets and data, see [User lifecycle and data handling in Amazon Quick](#).

Accounts can be deleted by either an AWS administrator or an Amazon Quick administrator. Deleting a Amazon Quick user account works the same in both the Standard and Enterprise editions of Amazon Quick.

Deleting a Amazon Quick user account removes or transfers their resources. In Enterprise edition, the network administrator can temporarily deactivate a Amazon Quick user account by removing it from the network group that has access to Amazon Quick. If a user is deleted, but not deactivated, that user can still access Amazon Quick as a new user. For more information about deactivating an Enterprise account, see [Deactivating user accounts](#).

Use the following procedure to delete a Amazon Quick user account.

1. Choose your user name on the application bar and then choose **Manage Quick**.
2. Choose **Manage Users**.
3. Locate the account you want to delete and then choose the delete icon at the end of that row.
4. Choose to either delete or transfer any resources owned by the user and then choose **OK**.
5. Do one of the following:

- If you chose to transfer user resources, enter the user name of the account to transfer them to and then choose **Delete and transfer resources**.
- If you chose to delete user resources, choose **Delete**. You can't undo this action.

User-driven license upgrades

Applies to: Enterprise Edition

Intended audience: System administrators and Amazon Quick administrators

User-driven license upgrades allow users to seamlessly transition between license tiers as their needs grow. When users encounter a feature that requires a higher tier or reach their usage limits, Amazon Quick provides contextual upgrade prompts to help them continue their work without interruption.

Note

Depending on your organization's settings, upgrades may process automatically or require administrator approval.

Overview

User-driven license upgrades enable Amazon Quick users to request higher-tier licenses when they encounter features or reach usage limits that require additional capabilities. This self-service approach reduces friction in the user experience while maintaining administrative control over licensing decisions.

The upgrade process is designed to be contextual and non-disruptive. Users see upgrade prompts at the moment they need additional capabilities, allowing them to continue their work seamlessly after approval.

Available upgrade paths

Users can request upgrades along the following paths:

- **BI Reader to Professional:** Enables advanced analytics, custom visualizations, and enhanced sharing capabilities.
- **Professional to Enterprise:** Unlocks enterprise features including advanced security, governance, and administrative controls.

Each upgrade path provides immediate access to new features upon approval, with billing changes taking effect according to your organization's billing cycle.

Feature-based upgrades

When users attempt to access features not available in their current tier, Amazon Quick presents contextual upgrade prompts. These prompts explain the benefits of the higher tier and guide users through the upgrade request process.

Common scenarios that trigger feature-based upgrade prompts include:

- Accessing advanced visualization types
- Attempting to use AI-powered features
- Trying to create or modify certain dashboard elements
- Accessing enterprise security or governance features

Threshold-based upgrades

Users receive upgrade prompts when they approach or exceed usage limits for their current license tier. These prompts appear in context, allowing users to upgrade and continue their work without interruption.

Threshold-based upgrades are triggered by:

- Reaching data source connection limits
- Exceeding dashboard or analysis creation limits
- Approaching user or sharing limits
- Hitting storage or compute usage thresholds

The system provides clear information about current usage, limits, and the benefits of upgrading to the next tier.

Topics

- [How users upgrade their license](#)
- [Administrator configuration](#)
- [Managing upgrade requests](#)
- [Monitoring upgrades](#)
- [User-driven license upgrades FAQ](#)

How users upgrade their license

Applies to: Enterprise Edition

To upgrade from the Home Page

1. Navigate to the Amazon Quick Home Page
2. Locate the "Explore the powerful new features of Amazon Quick" section at the top of the page
3. If needed, choose **Show More** to view all available feature cards
4. Select a feature card that interests you (such as AI features or Automation)
5. Review the feature description and benefits in the modal that appears
6. Choose **Request an upgrade to Professional** or **Request an upgrade to Enterprise**
7. Choose **Confirm Upgrade**

What happens next:

- **Auto-approval enabled:** Users see a message: "Your access is being upgraded. Check back in 10 min to continue." The upgrade processes automatically within approximately 10 minutes.
- **Admin approval required:** Users see a message: "Your request has been sent. An email will be sent when approved." The administrator will review and process the request.

To upgrade when reaching usage limits

- When users reach their usage threshold, a notification appears in their current workflow
- Users review their current usage and the limits of the next tier

- Choose **Upgrade Now** to continue work
- Review the pricing and features included in the new tier
- Choose **Confirm Upgrade**

To view and compare license tiers

- From any upgrade prompt, choose **View License Tiers**
- Review the features and pricing for each tier
- Compare current tier with available upgrade options
- Select the tier that best meets needs
- Choose **Upgrade** to proceed

Viewing current license

To check current license tier and usage:

1. Choose the profile icon in the top-right corner
2. The current license tier is displayed directly in the profile dropdown

Administrator configuration

Applies to: Enterprise Edition

Intended audience: System administrators and Amazon Quick administrators

As a Amazon Quick administrator, you can configure various settings for user-driven license upgrades, including enabling the feature, setting approval policies, and configuring custom permissions profiles.

Enabling user-driven upgrades

Before users can request license upgrades, you must enable the user-driven upgrade feature in your Amazon Quick account settings.

Access Administrator Settings

- Log in to your Amazon Quick administrator account
- Navigate to the 'Manage Account' section of the console
- Check for User upgrades Section
- Use the toggle button to Enable/Disable - Automatically approve upgrades

Configuring approval settings

Administrators can configure three types of upgrade governance:

- **Turned off altogether** — Users cannot request upgrades
- **Admin Approval** (default for IAM Identity Center and AD) — Users can request upgrades that require administrator approval
- **User Driven Upgrades** (default for IdP) — Users can upgrade automatically without approval

Creating a custom permissions profile for Amazon Quick account that is integrated with IAM Identity Center or Active Directory

Amazon Quick account administrators can use the following procedure to create a custom permissions profile for an Amazon Quick account that is integrated with IAM Identity Center or Active Directory.

To create a custom permissions profile for an Amazon Quick account that is integrated with IAM Identity Center or Active Directory

1. Sign in to the [AWS Management Console](#).
2. Open Amazon Quick.
3. The Amazon Quick Admin console opens. Choose **Custom Permissions**.
4. The **Manage custom permissions** page opens. Choose one of the following options.
5. To create a new custom permissions profile, choose **Create**.
6. To edit or view an existing custom permissions profile, choose the ellipses (three dots) next to the profile that you want, and then choose **Edit**.
7. If you want to create or update a custom permissions profile, make selections for the following items.
8. For **Name**, enter a name for the custom permissions profile.

9. For **Restrictions**, choose "Allow users to upgrade or request upgrades" under "User Upgrades". Choose any other options that you want to deny. Any option that you don't choose is allowed.
10. Choose **Create** or **Update** to confirm your choices. To go back without making any changes, choose **Back**.
11. Once you are done making changes, record the name of the custom permissions profile. Provide the name of the custom permissions profile to API users so that they can apply the custom permissions profile to roles or users.

Creating a custom permissions profile for an Amazon Quick account that uses Amazon Quick managed users

Amazon Quick account administrators can use the following procedure to create a custom permissions profile for an Amazon Quick account that uses Amazon Quick managed users.

To create a custom permissions profile for Amazon Quick managed users

1. Open the [Amazon Quick console](#).
2. From any page in the Amazon Quick console, choose **Manage Amazon Quick** at the top right corner.
3. Only Amazon Quick administrators have access to the **Manage Amazon Quick** menu option. If you don't have access to the **Manage Amazon Quick** menu, contact your Amazon Quick administrator for assistance.
4. Choose **Custom permissions**. You can also choose the **Manage users** section and then choose **Manage custom permissions**.
5. The **Manage custom permissions** page opens. Choose one of the following options.
6. To create a new custom permissions profile, choose **Create**.
7. To edit or view an existing custom permissions profile, choose the ellipses (three dots) next to the profile that you want, and then choose **Edit**.
8. If you want to create or update a custom permissions profile, make selections for the following items.
9. For **Name**, enter a name for the custom permissions profile.
10. For **Restrictions**, choose "Allow users to upgrade or request upgrades" under "User Upgrades". Choose any other options that you want to deny. Any option that you don't choose is allowed.
11. Choose **Create** or **Update** to confirm your choices. To go back without making any changes, choose **Back**.

12. Once you are done making changes, record the name of the custom permissions profile.
Provide the name of the custom permissions profile to API users so that they can apply the custom permissions profile to roles or users.

Applying custom permissions profiles

Apply a custom permissions profile to an Amazon Quick role with the Amazon Quick API

For more information, see [Apply custom permissions profile to role](#).

Apply a custom permissions profile to a user with the Amazon Quick API

For more information, see [Apply custom permissions profile to user](#).

Apply a custom permissions profile to an account

For more information, see [Apply custom permissions profile to account](#).

Apply a custom permissions profile to an account using the Amazon Quick APIs

For more information, see [Apply custom permissions profile using APIs](#).

Managing upgrade requests

Applies to: Enterprise Edition

Intended audience: System administrators and Amazon Quick administrators

As an administrator, you can view and manage user license upgrade requests through the Amazon Quick admin console. This section describes how to approve, deny, or verify upgrade requests based on your organization's identity management setup.

Viewing upgrade requests

To view pending upgrade requests, navigate to the User upgrades section in the admin console. The requests table displays the username, current role, requested role, request date, and available actions for each pending request.

To view upgrade requests

1. Sign in to the Amazon Quick console as an administrator.
2. Choose **Manage Amazon Quick** from the top right corner.
3. Navigate to the **User upgrades** section.
4. Review the pending requests in the table.

Processing upgrade requests

The process for handling upgrade requests varies depending on your organization's identity management setup.

Processing requests for IdP identity type

For organizations using Identity Provider (IdP) authentication, administrators can directly approve or deny requests.

To process IdP upgrade requests

1. In the upgrade requests table, choose **View details** for the request you want to process.
2. Review the request details in the modal that opens.
3. Choose one of the following options:
 - **Approve** — Grants the requested license upgrade immediately
 - **Deny** — Rejects the upgrade request
4. Confirm your selection. The system displays a success message and removes the processed request from the pending list.

Processing requests for IAM Identity Center and Active Directory identity types

For organizations using IAM Identity Center or Active Directory, administrators must add users to the appropriate groups before verifying the upgrade.

To process IAM Identity Center and Active Directory upgrade requests

1. In the upgrade requests table, choose **View details** for the request you want to process.
2. Note the required group name(s) displayed in the modal.
3. Navigate to your IAM Identity Center or Active Directory management console.

4. Add the user to the required group for the requested license tier.
5. Return to the Amazon Quick admin console and choose **Verify**.
6. The system checks the user's group membership and completes the upgrade if verification succeeds.

Alternatively, you can choose **Deny** to reject the upgrade request without making any group changes.

Understanding upgrade request statuses

Upgrade requests can have the following statuses:

- **Pending** — Request is awaiting administrator action
- **Approved** — Request has been approved and the user's license has been upgraded
- **Denied** — Request has been rejected by an administrator
- **Verified** — For IAM Identity Center/AD setups, the user has been added to the required group and the upgrade is complete

Monitoring upgrades

Applies to: Enterprise Edition

Intended audience: System administrators and Amazon Quick administrators

Administrators can monitor user upgrades through:

- **CloudTrail logs** — Provides audit trails of all upgrade requests and approvals.

Billing considerations for administrators

- Users are billed for their highest license tier within each billing period.
- **Prorated Billing:** When you upgrade a user subscription in Amazon QuickSuite, billing is prorated for the month of the upgrade. This means:

If you upgrade from BI Reader (\$3/month) to Professional (\$20/month), you're billed at the \$20/month rate from the time of upgrade until the end of that billing month.

If you then upgrade again from Professional (\$20/month) to Enterprise (\$40/month) within the same month, you're billed at the \$40/month rate from that second upgrade time until the end of the month.

User-driven license upgrades FAQ

Applies to: Enterprise Edition

What are user-driven license upgrades?

User-driven license upgrades allow Amazon Quick users to request license upgrades directly from the console, such as upgrading from Reader to Author or from Standard to Pro subscriptions. All requests require administrator approval.

Who can request license upgrades?

Any user with an active Amazon Quick account can request upgrades, provided the feature is enabled by their administrator and they meet the eligibility criteria for the requested upgrade type.

Can users downgrade their license tier?

License downgrades are not supported.

What happens to usage limits when upgrading?

Usage limits increase immediately to match the new tier. However, current usage does not reset—it carries over to the new tier.

How long does the approval process take?

- **Auto-approval:** Approximately 10 minutes
- **Admin approval:** Typically 1-2 business days, depending on administrator response time

Related topics

- [Managing subscriptions](#)

- [Managing user access](#)
- [Understanding license tiers](#)
- [Custom permissions](#)

Creating and managing groups in Amazon Quick

Intended audience: System administrators

Applies to: Enterprise Edition

Note

If your Amazon Quick account is integrated with IAM Identity Center (recommended), groups are not managed in the Quick application. Instead, groups are managed in IAM Identity Center or in the third-party identity provider that you configured in IAM Identity Center. Groups are synced automatically between Quick and IAM Identity Center.

Admins with IAM credentials who have access to the Quick console can organize sets of users into groups that make it easier to manage access and security. For example, you can create a group of users that you can share Amazon Quick assets with all at once. You can create and manage groups using the Amazon Quick console or the AWS Command Line Interface (AWS CLI). You can create up to 10,000 groups in a namespace. If you want to create more than 10,000 groups in a namespace, contact [AWS Support](#).

Use the topics below to create, and modify groups with in the Amazon Quick console or with the Amazon Quick APIs.

Topics

- [Create a group with the Amazon Quick console](#)
- [Change a group description with the Amazon Quick console](#)
- [Manage group membership in the Amazon Quick console](#)
- [Create and manage groups with the AWS CLI](#)

Create a group with the Amazon Quick console

Use the following procedure to create a group in the Quick console.

To create a user group in the Amazon Quick console:

1. On the Amazon Quick start page, choose **Manage Quick**, and then choose **Manage groups**.
2. Choose **NEW GROUP**.
3. On the **Create new group** page, enter the name and description of the new group in the corresponding boxes.
4. When you're finished, choose **Create** to create the new group.

Note

You can't delete a group from the Amazon Quick console, but you can delete a group with the AWS CLI. For more information on deleting a Amazon Quick group with the AWS CLI, see [Create and manage groups with the AWS CLI](#).

Change a group description with the Amazon Quick console

After you have created a new group, you can't change the group's title but you can change the group's description.

To change the description of a group:

1. On the Amazon Quick start page, choose **Manage Quick**, and then choose **Manage groups**.
2. Choose the group that you want to change, and then choose the **Edit** link next to the group description.
3. In the **Edit description** box that appears, enter the new description and choose **Save**.

Manage group membership in the Amazon Quick console

After you create a group, you can add and remove users from the **Manage groups** page. You can't add a user to a group if you haven't added the user to your account. For more information on adding users to your Amazon Quick account, see [Managing user access inside Amazon Quick](#).

To add a user to a group

1. On the Amazon Quick start page, choose **Manage Quick**, and then choose **Manage groups**.
2. Choose the group that you want to add a user to, and choose **ADD USER** at the page's upper right.
3. Enter the user name or email of the user that you want to add, and choose the correct user for **Search users**.

To remove a user from a group:

1. On the Amazon Quick start page, choose **Manage Quick**, and then choose **Manage groups**.
2. Choose the group that you want to remove a user from.
3. Find the user that you want to remove and choose **Remove**.

Choosing **remove** automatically removes the selected user from the group.

You can also search for a group member by entering the user's full user name into the search bar on the right-hand side of the group's page.

Create and manage groups with the AWS CLI

Before you begin, make sure that you have the AWS CLI installed. For more information, see [Installing the AWS CLI](#) in the *AWS CLI User Guide*.

Use the following procedure to create an Quick user group.

1. Open a terminal window. If you are using Microsoft Windows, open a command prompt.
2. Enter the following command at the prompt to create a group. Substitute the correct values for your parameters.

```
aws quicksight create-group --aws-account-id=111122223333 --namespace=default --  
group-name="Sales-Management" --description="Sales Management - Forecasting"
```

You might find it easier to create the command in a text editor before entering it at the prompt. For more information on `create-group` and other available commands, see the [Amazon Quick API reference](#).

3. Verify that the group exists by using a command similar to one of the following. The following command lists all groups.

```
aws quicksight list-groups --aws-account-id 111122223333 --namespace default
```

The following command describes a specific group.

```
aws quicksight describe-group --aws-account-id 111122223333 --namespace default --group-name Sales
```

The following command searches for groups in a specified Amazon Quick namespace.

```
aws quicksight search-groups --region us-west-2 --aws-account-id 111122223333 --namespace default --filters "[{\"Operator\": \"StartsWith\", \"Name\": \"GROUP_NAME\", \"Value\": \"Mar\"}]"
```

4. Add a member to the new group by using a command similar to the following.

```
aws quicksight create-group-membership --aws-account-id 111122223333 --namespace default --group-name Sales --member-name Pat
```

The following command determines if a user is a member of a specified group.

```
aws quicksight describe-group-membership --region us-west-2 --aws-account-id 111122223333 --namespace default --group-name Marketing-East --member-name user
```

Enter the following command at the prompt to delete a group. Substitute the correct values for your parameters.

```
aws quicksight delete-group --aws-account-id 111122223333 --namespace default --group-name Marketing-East
```

You might find it easier to create the command in a text editor before entering it at the prompt. For more information on `delete-group` and other available commands, see the [Amazon Quick API reference](#).

Managing security

Security management in Amazon Quick encompasses multiple layers of protection to ensure your data and resources remain secure. This section covers key security features including domain

management for controlling access to specific organizational domains, namespace configuration for multi-tenant environments, mobile device security settings, IP address and VPC restrictions, VPC connection management, and customer-managed encryption keys. These security controls help you implement defense-in-depth strategies and meet your organization's compliance requirements while maintaining the flexibility to share insights securely across your organization.

Topics

- [Managing domains](#)
- [Supporting multitenancy with isolated namespaces](#)
- [Managing mobile settings](#)
- [Managing IP/VPC restrictions](#)
- [Managing VPC connection in Amazon Quick](#)
- [AI guardrails in Amazon Quick](#)
- [Encrypting your Amazon Quick data with AWS Key Management Service customer managed keys](#)

Managing domains

Applies to: Enterprise Edition

Intended audience: Amazon Quick administrators

In Amazon Quick Enterprise edition, you can embed Amazon Quick Sight dashboards, visuals, consoles, and Q search bars in an app or web page. Domains that are going to host these embedded assets must be on an *allow list*, the list of approved domains for your Quick subscription. This requirement protects your data by preventing unapproved domains from hosting embedded dashboards. To embed a Amazon Quick Sight dashboard, visual, console, or Q search bar to a web page or app, add approved domains to a static allow list in the Quick console. Alternatively, add them at runtime with the Quick API.

Use the following sections to learn more about adding domains for embedded analytics.

Topics

- [Allow listing static domains](#)

- [Allow listing domains at runtime with the Amazon Quick API](#)

Allow listing static domains

You can add static domains to your allow list through the Amazon Quick console. All domains on your allow list (such as development, staging, and production) must be explicitly allowed, and they must use HTTPS. You can add up to 100 domains to the allow list.

To embed a dashboard to a static domain:

- Approve the hosting domains and subdomains for embedding.
- Publish the dashboard.
- Share the dashboard with users or groups so they can see the embedded version of it.

Use the following procedure to view or edit the list of approved domains.

To view or edit the list of approved domains

1. Choose the profile icon at top right.
2. Choose **Manage Amazon Quick**. You must be an Amazon Quick admin to access this screen.
3. Choose **Domains and Embedding** on the left. The domains that you can embed a dashboard in are listed at the bottom of the page.
4. (Optional) Add a new domain here by entering it in the **Domain** box. You can also choose **Include subdomains** to allow embedded dashboards on all subdomains. Choose **Add** to add the domain.

You can edit or delete existing domain by choosing the icons next to each domain in the list at the bottom of the page.

Make sure that you use a valid HTTPS URL. The following list shows examples of URLs that are valid for embedded dashboards that use a static domain:

- https://example-1.com
- https://www.アマゾンドメイン.jp
- https://www.亚马逊域名.cn:1234
- https://111.222.33.44:1234

- `https://111.222.33.44`
- `http://localhost`

The following list shows examples of URLs that are *not* valid for embedded dashboards:

- `http://example`
- `https://example.com.*.example-1.co.uk`
- `https://co.uk`
- `https://111.222.33.44.55:1234`
- `https://111.222.33.44.55`

Allow listing domains at runtime with the Amazon Quick API

You can add a domain at runtime to an allow list with the `AllowedDomains` parameter of a `GenerateEmbedUrlForAnonymousUser` or a `GenerateEmbedUrlForRegisteredUser` API call. The `AllowedDomains` parameter is an optional parameter. It grants you the option as a developer to override the static domains that are configured in the **Manage Amazon Quick** menu.

You can list up to three domains or subdomains. Adding domains to the allow list at runtime also adds HTTP support for the domain `localhost`. The generated URL is then embedded in a developer's website. Only the domains that are listed in the parameter can access the embedded dashboard.

Security best practice for IAM condition operators

Improperly configured IAM condition operators can allow unauthorized access to your embedded Quick resources through URL variations. When using the `quicksight:AllowedEmbeddingDomains` condition key in your IAM policies, use condition operators that either allow specific domains or deny all domains that are not specifically allowed. For more information about IAM condition operators, see [IAM JSON policy elements: Condition operators](#) in the IAM User Guide.

Many different URL variations can point to the same resource. For example, the following URLs all resolve to the same content:

- `https://example.com`
- `https://example.com/`

- `https://Example.com`

If your policy uses operators that do not account for these URL variations, an attacker can bypass your restrictions by providing equivalent URL variations.

You must validate that your IAM policy uses appropriate condition operators to prevent bypass vulnerabilities and ensure that only your intended domains can access your embedded resources.

To embed a dashboard to a domain at runtime, see [Embedding with the Amazon Quick APIs](#).

Make sure that you use a valid URL. The following list shows examples of URLs that are valid for embedded dashboards that use a runtime domain:

- `https://example-1.com`
- `http://localhost`
- `https://www.アマゾンドメイン.jp`
- `https://*.sapp.amazon.com`

The following list shows examples of URLs that are *not* valid for embedded dashboards:

- `https://example.com.*.example-1.co.uk`
- `https://co.uk`
- `https://111.222.33.44.55:1234`
- `https://111.222.33.44.55`

For more information about embedded dashboards, see [Embedding with the Amazon Quick APIs](#).

Supporting multitenancy with isolated namespaces

Amazon Quick Enterprise edition supports multitenancy through namespaces. An Amazon Quick *namespace* is a logical container that you can use to organize clients, subsidiaries, teams, and so on. Namespaces can help you achieve the following goals:

- Your Amazon Quick subscription allows users to discover shared content and share with other users. At the same time, in the Amazon Quick end-user console experience, users in

one namespace can't see or interact with users in another namespace. The Amazon Quick administrative console uses a different permission model. For more information, see the following important note.

- You can securely isolate data and also support diverse workloads without adding additional AWS accounts. Access to data is still strictly controlled by AWS security features. Users can see assets (like data and dashboards) only if they have the correct resource permissions. Also, users who have permissions can't inadvertently expose content to people who outside of their namespace. For more information, see [Supporting multitenancy with isolated namespaces](#).
- You can monitor data flows and usage reports, neatly partitioned by namespace. Categorizing data and reports by namespace can help simplify cost and security analysis.
- After you've registered users into your namespace, there's no additional administrative complexity or overhead.
- Namespaces are designed to span AWS Regions, so the use containment doesn't change even if a person signs in to a different AWS Region.

Namespace isolation and the administrative console

Amazon Quick uses two different permission models depending on the console experience. In the Amazon Quick end-user console, namespace-level isolation is enforced. Each user's Amazon Quick role determines the actions that they can perform, and their namespace determines which users and groups they can interact with.

In the Amazon Quick administrative console, only AWS Identity and Access Management (IAM) permissions are enforced. Users operate outside of the namespace context, and the user's IAM policies determine access, not their namespace membership.

Users can operate under both permission models. When they use the end-user console, their Amazon Quick role and namespace determine their access. When they use the administrative console, their IAM permissions determine their access. For example, an IAM administrator with the appropriate permissions can view and manage users and groups across namespaces from the administrative console.

Namespaces currently have the following limitations:

- Custom namespaces—those that are not the default namespace—are only accessible to IAM Federated Single-Sign On users.
- Use default namespaces instead of custom namespaces if you need to support the following:

- Integrating your Amazon Quick account with IAM Identity Center. For more information on integrating your Amazon Quick account with IAM Identity Center, see [AWS security in Amazon Quick](#).
- Password-based logins.
- Credential-based Active Directory logins.
- You can't transfer users directly from one namespace to another. You can choose to do some or all of this work programmatically. For more information, see the [Quick API reference](#). At the bottom of the page of each API operation, there's a list of links to the same operation in the SDKs for other languages. To see what SDKs are available, see [SDKs and toolkits](#) in the [AWS getting started resource center](#).
- Namespaces are useful for isolating users and permissions, but not for sharing assets. Dashboards, datasets, and analyses can be shared with users in different namespaces. By default, users can't access items that exist in the same namespace by default, but gain access to specific assets when the asset is shared with them.

If you don't have an existing AWS account or you need to sign up for Amazon Quick, read the following guidelines, then follow the applicable instructions in [Signing up for an Amazon Quick subscription](#):

- Sign up for Enterprise edition.
- When asked which method you want to connect with, choose **Role Based Federation (IAM)**. Currently, namespaces support only customers who use an AWS Identity and Access Management (IAM) role with a web identity federation. For more information, see [Creating a role for a third-party Identity Provider \(federation\)](#)
- Complete the process of signing up.
- Use the Amazon Quick [CreateNamespace](#) API operation to create one or more namespaces.
- To start adding users, first follow the instructions in [Setting up IdP federation using IAM and Amazon Quick](#). Then use the [RegisterUser](#) API operation to add users to the appropriate namespace.

If you already signed up for Standard edition, you can easily upgrade your subscription to Enterprise edition. The person performing the upgrade must be a Amazon Quick user with administrator privileges. For more information, see [Upgrading your Amazon Quick subscription](#).

If you have an Enterprise edition subscription that you've been using for some time, it's also possible to migrate your users into namespaces. When you sign up for Amazon Quick and add users, all of them reside in the default namespace. All of the users can interact directly with each other and share data and dashboards with each other. To isolate your users from each other, you can create one or more additional namespaces.

 Important

Amazon Quick assets and resources, including datasets, data sources, dashboards, analyses, and so on, exist outside of any namespace. They're visible only to users who have resource permissions granted to them.

To implement namespaces, you use the following Amazon Quick API operations:

- [CreateNamespace](#)
- [DescribeNamespace](#)
- [ListNamespaces](#)
- [DeleteNamespace](#)

Namespaces are not supported in the Regions listed below:

- af-south-1 Africa (Cape Town)
- ap-southeast-3 Asia Pacific (Jakarta)
- ap-southeast-5 Asia Pacific (Malaysia)
- eu-south-1 Europe (Milan)
- eu-central-2 Europe (Zurich)

 Note

If you need to install the AWS CLI, see [Installing the AWS CLI version 2](#) in the *AWS Command Line Interface User Guide*.

To add users to a namespace, you use the [RegisterUser](#) API operation. Each namespace has a completely independent set of users. The user ARNs include the namespace qualifier to distinguish them, as shown in the following examples:

- Amazon Quick considers these two entities to be different persons:
 - `arn:aws:quicksight:us-east-1:111122223333:user/namespace-123/username123`
 - `arn:aws:quicksight:us-east-1:111122223333:user/namespace-456/username123`
- Amazon Quick considers these two entities to be the same person:
 - `arn:aws:quicksight:us-east-1:111122223333:user/namespace-123/username123`
 - `arn:aws:quicksight:us-west-2:111122223333:user/namespace-123/username123`

When you use [RegisterUser](#), you select an access level for each user. After a person's user name is assigned to one of the security cohorts, their access to the console and API is restricted. People using Amazon Quick can have a single access level, as follows:

- Reader access, for read-only subscribers of a dashboard
- Author access, for analysts and dashboard designers
- Admin access, for Amazon Quick administrators

To migrate existing users in one namespace to a different namespace

Follow the procedure below to migrate existing users from one namespace to a different namespace.

1. Identify the users that you want to transfer to a different namespace by using the Amazon Quick user and group API operations. For more information, see [API operations for controlling access](#) in the [Quick API reference](#).
2. Create users in the new namespace by using the [RegisterUser](#) API operation. Within a namespace, user names are unique.

If a namespace user starts using the Amazon Quick console or API in a new AWS Region, that user is still constrained to the namespace that you added them to. Each namespace represents

a user directory of an identity provider. As such, it originates in the primary AWS Region where Amazon Quick is set up. However, because the user directory is propagated globally in your AWS account, the namespace is accessible from any AWS Region where your users are using Amazon Quick.

3. To identify the asset and resource permissions that the new namespace users need, use the Amazon Quick API operations associated with each type of asset (dashboards, datasets, and so on). For more information, see [QuickSight API operations to control assets](#) in the [Quick API reference](#).

For example, let's say you are focusing on dashboards. You can use `ListDashboards` to list all the dashboard IDs in your AWS account. Then, to determine which users or groups can access these dashboards, you can use `DescribeDashboardPermissions` on the result set generated by `ListDashboards`. If you need to identify specific versions of a dashboard, you can `ListDashboardVersions` for that. You can also collect information about the location of the data that's used in the dashboard with the data source and dataset API operations. For more information, see [QuickSight API operations to control data resources](#) in the [Quick API reference](#).

For more information about filtering API response output, see the SDK documentation for the language you're using. For information relating to the AWS Command Line Interface (AWS CLI), see [Controlling command output from the AWS CLI](#) in the [AWS Command Line Interface User Guide](#).

4. For Amazon Quick assets and resources, copy the permissions that the source namespace user has for each asset. Then use, for example, `UpdateDashboardPermissions` to apply the same permissions to the target namespace user. Each asset type has its own separate set of API operations for controlling the permissions that users have to use it. For more information, see [QuickSight API operations for asset and resource permissions](#) in the [Quick API reference](#).
5. When you are finished adding users and permissions, it's a good practice to allow some time for user acceptance testing. Doing this ensures that everyone is successfully using the new namespace. It also ensures that all assets and resources are accessible in the new namespace.

After you're certain that you no longer need the original user names, you can begin to deprecate their permissions in the original namespace. Finally, when the users are ready, you can remove the unused group and user names in the source namespace. Do this in each AWS Region where your users were previously active.

Managing mobile settings

Mobile settings allow Amazon Quick admins to enhance data security and maintain compliance standards for users accessing Quick through mobile devices. These settings apply to all users accessing Quick through the mobile application.

To access mobile settings:

1. Open the [Quick console](#).
2. Choose the user icon at the top right, and then choose **Manage Amazon Quick**.
3. Under the **Security** section, choose **Mobile settings**.
4. Choose settings to enhance your data security on mobile devices:
 - **Require use of biometrics or PIN to unlock device** - Check the box to enable this security requirement. When enabled, users must use fingerprint recognition, Face ID, device PINs, or other biometric authentication methods supported by their device to access Amazon Quick.
 - **Require devices to use the latest version of the mobile operating system** - Check the box to enable this security requirement. When enabled, users must keep their mobile devices' operating systems up to date to access Amazon Quick.
5. To manage how long users can remain logged in on mobile devices:
 - a. Locate the **Current session lifetime for all devices** setting.
 - b. Select the desired duration from the dropdown menu. The maximum session lifetime depends on your identity provider:
 - **IAM Identity Center**: Maximum session lifetime is 7 days.
 - **All other identity providers**: Maximum session lifetime is 30 days (default).

The minimum session lifetime for all identity providers is 12 hours.

- c. Choose **Update** to apply your changes.

Best practices

- Regularly review and update mobile security settings to align with your organization's security policies
- Set an appropriate session lifetime that balances security with user convenience

- Enable biometric/PIN requirements for additional security, especially when dealing with sensitive data
- Keep operating system requirements enabled to ensure devices have the latest security updates
- Communicate changes in mobile settings to your users before implementation

Managing IP/VPC restrictions

Turning on Internet Protocol (IP) and VPC endpoint restrictions in Amazon Quick

Applies to: Enterprise Edition

You can limit access to your organization's Amazon Quick account to a predefined list of IP ranges, VPC IDs, and VPC endpoint IDs. For example, you can create an IP rule that allows users to access your Quick account only from IP addresses associated with your company's office or remote virtual private network (VPN). You can also create a VPC endpoint rule that allows users access to your Quick account only from the VPC that is used for Direct Connect.

For more information about setting up VPC endpoints in Quick, see [Quick Interface VPC endpoints \(AWS PrivateLink\)](#) for more information about how to setup VPC endpoints.

Note

For Quick accounts integrated with IAM Identity Center, browser-based authentication requires access to the public AWS sign-in and IAM Identity Center endpoints. VPC endpoints do not provide this browser sign-in path. You don't need public IP addresses. Outbound internet connectivity, such as a network address translation (NAT) gateway or proxy, can provide access to these endpoints. For more information, see [Enable access to the IAM Identity Center access portal](#) and [Considerations for private access](#) in the *IAM Identity Center User Guide*.

Only admins with AWS Identity and Access Management (IAM) credentials who have access to the Quick console pages can access the IP and VPC endpoint restrictions table.

Topics

- [Adding an IP or VPC endpoint rule](#)

- [Update an existing rule](#)
- [Delete a rule](#)
- [Turning on your IP and VPC endpoint rules](#)

Adding an IP or VPC endpoint rule

An *IP rule* is created when you add a CIDR address with a public IP version 4 address to the restrictions table. A *VPC endpoint rule* is created when you add either a VPC ID or a VPC endpoint ID to the restrictions table. You can add up to 100 IP and VPC endpoint rules combined to the restrictions table. To request a higher limit, contact AWS Support. You can only add rules from the AWS Region where your account is. All traffic that is not allowed by either the IP rule or the VPC endpoint rule is blocked when the restriction is turned on.

A *CIDR address* is composed of two parts: the prefix and the suffix. The prefix is the CIDR's network address and is written like a normal IP address. The suffix shows how many bits are in the address. An example of a complete CIDR address is `10.24.34.0/23`.

IP and VPC endpoint rules apply only to Quick web, embedded, and mobile access and don't restrict access to the public API. Your users can still call all API operations from restricted IP ranges. For information on restricting calls to the public API from specific IP addresses, see [AWS: Denies access to AWS based on the source IP](#) in the *IAM User Guide*.

Before you save any rule changes or turn on other rules, make sure that you have a rule that includes your IP address or VPC endpoint ID. If there isn't a rule that allows your traffic, you can't save your changes.

To add an IP or VPC endpoint rule

1. On the Quick start page, choose **Manage QuickSight**, and then choose **Security and Permissions**.
2. Choose **IP and VPC endpoint restrictions**.
3. Perform one of the following actions.
 - a. For **IP restriction**, enter the CIDR address that defines the IP range that you want to create a rule for.
 - b. For **VPC endpoint restriction**, enter either the VPC ID or the VPC endpoint ID of the endpoint that you want to create a rule for.

4. (Optional) For **Description**, enter a description for the rule. Doing this can help you differentiate your rules.
5. Choose **Add**.
6. Choose **Save changes** in the box that appears to apply the rule.

It can take up to 10 minutes for a rule to be fully implemented.

Update an existing rule

After you create an IP or VPC endpoint rule, use the IP and VPC restrictions table to make changes to the rule. Use the following procedure to update an existing IP or VPC endpoint rule in the IP and VPC restrictions table.

To update an existing IP or VPC endpoint rule

1. On the Quick start page, choose **Manage QuickSight**, and then choose **Security and Permissions**.
2. Choose **IP and VPC endpoint restrictions**.
3. Choose the edit icon to the right of the rule that you want to change.
4. Make your changes and choose **Update**.
5. Choose **Save changes** in the box that appears to update the rule.

It can take up to 10 minutes for an updated rule to be fully implemented.

Delete a rule

Use the following procedure to delete an IP or VPC endpoint rule from the IP and VPC endpoint restrictions table.

To delete an IP or VPC endpoint rule

1. On the Quick start page, choose **Manage QuickSight**, and then choose **Security and Permissions**.
2. Choose **IP and VPC endpoint restrictions**.
3. Locate the rule that you want to delete, and then choose the delete icon next to it. The rule appears with a strikethrough to show that it is marked for deletion.
4. Choose **Save changes** in the box that appears to delete the rule.

It can take up to 10 minutes for an updated rule to be deleted.

Turning on your IP and VPC endpoint rules

You can turn on or turn off your account's IP and VPC endpoint restrictions by using the **Rules** option at the top of the IP and VPC restrictions page. When rules are turned on, users from sources that are not on the restrictions table can't access Quick mobile, embedded, and website pages. IP and VPC endpoint rules are global and apply to all AWS Regions.

If a user is accessing the Quick account from a source that is not on the rule list when you turn on restrictions, they lose access to the account.

Account holders can audit users who make changes to the IP and VPC endpoint restrictions table by using AWS CloudTrail. For more information, see the [AWS CloudTrail User Guide](#).

Managing VPC connection in Amazon Quick

Applies to: Enterprise Edition

Intended audience: System administrators and Amazon Quick administrators

With Amazon Quick Enterprise Edition, account admins can configure a secure, private VPC connection to a Amazon Quick account from the Amazon Quick console or from the Amazon Quick CLI. Read the following walkthroughs on how to create, edit, and delete a VPC connection from a Amazon Quick account.

Topics

- [Configuring the VPC connection in the Amazon Quick console](#)
- [Configuring the VPC connection with the Amazon Quick CLI](#)
- [Testing the connection to your VPC data source](#)

Configuring the VPC connection in the Amazon Quick console

To create a secure private connection to the Amazon VPC service from the Amazon Quick console, use the following procedure.

Prerequisites

- Sign in to Amazon Quick as a Amazon Quick admin to set up a VPC connection in Amazon Quick. To verify that you're a Amazon Quick administrator, choose your profile icon in the upper-right. If your profile menu contains the option **Manage Amazon Quick**, then you're a Amazon Quick administrator. Make sure your admin role in IAM includes the following permissions. The "iam:PassRole" permission needs to be applied only to the execution role that is created in the procedures below.
 - "quicksight:ListVPCConnections"
 - "quicksight:CreateVPCConnection"
 - "quicksight:DescribeVPCConnection"
 - "quicksight>DeleteVPCConnection"
 - "quicksight:UpdateVPCConnection"
 - "ec2:describeSubnets"
 - "ec2:describeVpcs"
 - "ec2:describeSecurityGroups"
 - "iam:ListRoles"
 - "iam:PassRole"

The following example shows an IAM policy that applies "iam:PassRole" only to the execution role.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "iam:PassRole"
      ],
      "Resource": "arn:aws:iam::111122223333:role/vpc-role-for-qs"
    }
  ]
}
```

- Before you begin, make sure that you have the following information available to copy and paste into the **VPC Connection** screen. For more information, see [Finding information to connect to a VPC](#).
 - AWS Region – The AWS Region where you plan to create a connection to your data source.
 - VPC ID – The ID of the VPC that contains the data, the subnets, and the security groups that you plan to use.
 - Execution role– An IAM role that contains a trust policy that allows Amazon Quick to create, update, and delete network infrastructure in your account. This policy is required for all VPC connections. At minimum, the IAM policy needs the following Amazon EC2 permissions:
 - DescribeSecurityGroups
 - DescribeSubnets
 - CreateNetworkInterface
 - DeleteNetworkInterface
 - ModifyNetworkInterfaceAttribute

The following example shows an IAM policy that you can add to an existing IAM role to create, delete, or modify a VPC connection:

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ec2:CreateNetworkInterface",
        "ec2:ModifyNetworkInterfaceAttribute",
        "ec2>DeleteNetworkInterface",
        "ec2:DescribeSubnets",
        "ec2:DescribeSecurityGroups"
      ],
      "Resource": "*"
    }
  ]
}
```

After you add the necessary permissions to an IAM role, attach a trust policy to allow Amazon Quick to configure the VPC connection to your account. The following example shows a trust policy that you can add to an existing IAM role to allow Amazon Quick access to the role:

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": "quicksight.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

- Subnet IDs – The IDs of the subnets that the Amazon Quick network interface is using. Each VPC connection needs at least two subnets.
- Security group IDs – The IDs of the security groups. Each VPC connection needs at least one security group.

To create a secure private connection to the Amazon VPC service from Quick Enterprise edition

1. In Amazon Quick, choose your profile icon in the upper-right, then choose **Manage Amazon Quick**.

Only Amazon Quick administrators can view the **Manage Amazon Quick** option. If you don't see this option on your profile menu, you're not an administrator. In this case, contact your Amazon Quick account administrators for assistance.

2. In the left navigation pane, choose **Manage VPC connections**.
3. On the **Manage VPC connections** page that opens, choose **Add VPC connection**.
4. For **VPC connection name**, enter a unique descriptive name of your choice. This name doesn't need to be an actual VPC ID or name.
5. In the **VPC ID** dropdown menu, choose the ID of the VPC in Amazon EC2 that you want to connect to your Amazon Quick account. This field can't be changed later.

6. In the **Execution role** dropdown menu, choose the appropriate IAM role to use for the VPC connection. The **Execution role** dropdown only shows IAM policies that contain a trust policy that allows Amazon Quick to configure the VPC connection to your account.
7. In the **Subnets** table, choose a subnet ID from the **Subnet ID** dropdown menu of at least two of the listed **Availability zones**. The Availability Zones listed in the **Subnets** table are determined based on how you configured the VPC connection in the Amazon EC2 console.
8. (Optional) If you aren't using DNS resolver endpoints, skip to the next step.

If your database host IP address must be resolved through private DNS servers in your AWS account, enter IP addresses for Route 53 Resolver inbound endpoints (one per line).

Make sure that you are entering an endpoint, rather than a database address like the one you plan to use in Amazon Quick. Most databases that are hosted by AWS don't need to resolve DNS queries between VPCs and a customer's network. For more information, see [Resolving DNS queries between VPCs and your network](#) in the *Amazon Route 53 Developer Guide*. You only need this endpoint if you can't resolve the IP address that connects to your database by using the public DNS server system.

9. Review your choices, then choose **ADD**.

When you finish creating a VPC connection, the new connection appears in the **Manage VPC connections** table. In some cases, the status of the new VPC might be **UNAVAILABLE** until the connection is configured on the backend. After Amazon Quick is finished configuring the new connection, the status of the connection switches to **AVAILABLE**, which indicates that the connection has been established. The following table describes the different **Status** values for a VPC connection.

Status	Description
AVAILABLE	The VPC connection is established and can be used.
PARTIALLY AVAILABLE	One of the network interfaces that is configured to the VPC connection is unavailable. The VPC connection can still be used.
UNAVAILABLE	The VPC connection is not established and can't be used.

To see a summary of a VPC connection, choose a VPC connection from the **VPC connection name** row of the **Manage VPC connections** table. The pop-up box that appears shows information about the network interfaces associated with the VPC connection.

The following table describes the different **Status** values for a network interface.

Status	Description
CREATING	The network interface creation is in progress.
AVAILABLE	The network interface is available for use.
CREATION_FAILURE	The network interface couldn't be created.
UPDATING	The security group associated with the network interface is updating.
UPDATE_FAILED	The security group associated with the network interface did not update successfully.
DELETING	The network interface is in the process of being deleted.
DELETED	The network interface is deleted and can no longer be used.
DELETION_FAILED	The network interface deletion failed and can still be used.
DELETION_SCHEDULED	This network interface is scheduled for deletion.
ATTACHMENT_FAILED_ROLLBACK_FAILED	The elastic interface failed to attach and Amazon Quick was unable to delete the elastic network interface that was created within your account.

When you delete a network interface from a VPC connection, the status of the connection changes to **PARTIALLY AVAILABLE** to indicate the loss of a network interface.

To make changes to an existing VPC connection, choose the more actions (three-dots) button to the right of the connection that you want to modify, and choose **Edit**. In the **Edit VPC connection** window that appears, make your changes, and then choose **SAVE**.

To delete a VPC connection, choose the more actions (three-dots) button to the right of the connection that you want to delete and choose **Delete**. In the **Delete Amazon Quick VPC Connection** pop-up that appears, confirm that you want to delete the connection, and then choose **Delete**.

Configuring the VPC connection with the Amazon Quick CLI

To create a secure private connection to the Amazon VPC service from Quick with the Amazon Quick CLI, use the following procedure:

Prerequisites

- Before you begin, make sure that you have the following information available to copy and paste into the **VPC Connection** page. For more information, see [Finding information to connect to a VPC](#).
 - AWS Region – The AWS Region where you plan to create a connection to your data source.
 - VPC ID – The ID of the VPC that contains the data, the subnets, and the security groups that you plan to use.
 - Execution role– An IAM role that contains a trust policy that allows Amazon Quick to create, update, and delete network infrastructure in your account. This policy is required for all VPC connections. At minimum, the IAM policy needs the following Amazon EC2 permissions:
 - DescribeSecurityGroups
 - DescribeSubnets
 - CreateNetworkInterface
 - DeleteNetworkInterface
 - ModifyNetworkInterfaceAttribute

The following example shows an IAM policy that you can add to an existing IAM role to create, delete, or modify a VPC connection:

JSON

```
{  
  "Version": "2012-10-17",
```

```

    "Statement": [
      {
        "Effect": "Allow",
        "Action": [
          "ec2:CreateNetworkInterface",
          "ec2:ModifyNetworkInterfaceAttribute",
          "ec2>DeleteNetworkInterface",
          "ec2:DescribeSubnets",
          "ec2:DescribeSecurityGroups"
        ],
        "Resource": "*"
      }
    ]
  }
}

```

After you add the necessary permissions to an IAM role, attach a trust policy to allow Amazon Quick to configure the VPC connection to your account. The following shows an example trust policy that you can add to an existing IAM role to allow Amazon Quick access to the role:

JSON

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": "quicksight.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}

```

- Subnet IDs – The IDs of the subnets that the Amazon Quick network interface is using. Each VPC connection needs at least two subnets.
- Security group IDs – The IDs of the security groups. Each VPC connection needs at least one security group.

Using the AWS CLI

The following example creates a VPC connection.

```
aws quicksight create-vpc-connection \  
--aws-account-id 123456789012\  
--vpc-connection-id test \  
--name test \  
--subnet-ids ['subnet-12345678', 'subnet-12345678'] \  
--security-group-ids ['sg-12345678'] \  
--role-arn arn:aws:iam::123456789012:role/test-role \  
--region us-west-2
```

After you create a VPC connection, you can update, delete, or request a summary of the VPC connection.

The following example updates a VPC connection.

```
aws quicksight update-vpc-connection \  
--aws-account-id 123456789012 \  
--vpc-connection-id test \  
--name test \  
--subnet-ids ['subnet-12345678', 'subnet-12345678'] \  
--security-group-ids ['sg-12345678'] \  
--role-arn arn:aws:iam::123456789012:role/test-role \  
--region us-west-2
```

The following example deletes a VPC connection.

```
aws quicksight delete-vpc-connection \  
--aws-account-id 123456789012 \  
--vpc-connection-id test \  
--region us-west-2
```

The following example describes a VPC connection.

```
aws quicksight describe-vpc-connection \  
--aws-account-id 123456789012 \  
--vpc-connection-id test \  
--region us-west-2
```

The following table describes the different **Status** values for a network interface that `describe-vpc-connection` returns.

Status	Description
CREATING	The network interface creation is in progress.
AVAILABLE	The network interface is available for use.
CREATION_FAILURE	The network interface couldn't be created.
UPDATING	The security group associated with the network interface is updating.
UPDATE_FAILED	The security group associated with the network interface did not update successfully.
DELETING	The network interface is in the process of being deleted.
DELETED	The network interface is deleted and can no longer be used.
DELETION_FAILED	The network interface deletion failed and can still be used.
DELETION_SCHEDULED	This network interface is scheduled for deletion.
ATTACHMENT_FAILED_ROLLBACK_FAILED	The elastic interface failed to attach and Amazon Quick was unable to delete the elastic network interface that was created within your account.

You can also use the AWS CLI to generate a list of all VPC connections in your Amazon Quick account.

```
aws quicksight list-vpc-connections \
--aws-account-id 123456789012 \
```

```
--region us-west-2
```

Testing the connection to your VPC data source

Applies to: Enterprise Edition

Intended audience: Amazon Quick administrators and authors

To test whether you can connect to your data source through an existing Quick VPC connection, use the following procedure.

Before you begin, collect the information you need to connect. If you plan to copy and paste settings from a file, make sure that it doesn't contain any of the following: formatting (list bullets or numbers), blank space (spaces, tabs), or invisible "gremlin" (non-ASCII, null (ASCII 0), or control) characters.

1. On the Amazon Quick start page, choose **Manage data**.
2. On the **Datasets** page, choose **New data set**.
3. In the **FROM NEW DATA SOURCES** section of the **Create a data set** page, choose a supported data source that you want to connect to. For a list of data sources that support VPC, see [Identify the data sources to use](#).

Your data source instance must use the same VPC that you used to create the VPC connection. Also, the associated security group must be properly configured. For more information, see [Setting up a VPC to use with Amazon Quick](#).

4. Enter the connection information for the data source. The fields for the data source are sometimes displayed in different order depending on which data source you choose. For more information, see [Creating a data source](#).
 - For **Data source name**, enter a descriptive name for the new data source. This name appears beside the data source logo on a tile on the **Create a data set** page. For testing purposes, name it "**VPC test -**" followed by the database name or location, whichever is unique.
 - For **Connection Type**, choose the name of the VPC connection that has a route to your data source. If the correct VPC is missing from the list, ask a Amazon Quick administrator to

verify that the VPC connection is correct in Amazon Quick. If it looks correct, ask a system administrator to verify that the data source and VPC are set up for this purpose.

- The name or other identifier for the server or instance to connect to. The descriptors vary depending on which one you're connecting to, but it's usually one or more of the following: hostname, IP address, cluster ID, instance ID, connector, or site based URL.
- **Database name** shows the default database for the **Instance ID** cluster or instance. If you want to use a different database on that cluster or instance, enter its name.
- The name of the collection of data that you want to use.

The descriptor varies depending on the provider, but it's usually one of the following: database, warehouse, or catalog. In this topic, we use the word "database" as a generic term.

- For **Credentials**, enter a username and password to use for everyone who connects from Amazon Quick using this data source. The username must have permissions to do the following:
 - Access the target database.
 - Read (perform a SELECT statement on) all of the tables that you want to use in that database.
5. Choose **Validate connection** to verify your connection information is correct. If your connection doesn't validate, correct the connection information and try again. If the information looks correct but doesn't validate, do one or all of the following:
- Contact your data source administrator to verify your connection settings.
 - Contact your Amazon Quick administrator to verify the settings in the Amazon Quick VPC connection.
 - Contact your AWS administrator to verify that the VPC is correctly configured for use with Amazon Quick.
6. After the connection validates, choose **Create data source** to save the connection profile. Or, choose **Cancel** if you don't need to save it (recommended) after testing is complete.

AI guardrails in Amazon Quick

Amazon Quick takes a defense-in-depth approach to help keep chat interactions safe and responsible in the Amazon Quick web experience. Multiple safety layers work together, each addressing different risks at different points in an interaction. No single layer is designed to catch every risk on its own.

Usage of Amazon Quick AI features is subject to the [AWS Responsible AI Policy](#) on the AWS website.

The following topics describe how these controls work and how to configure them.

Topics

- [How AI safety controls work together](#)
- [Blocked words and phrases](#)

How AI safety controls work together

Amazon Quick uses the following controls together:

- **Built-in safety screening** – Amazon Bedrock Guardrails and Quick safeguards screen user requests for harmful content. This includes hate speech, insults, sexual content, violence, and misconduct. These checks also assess user requests for prompt attacks.
- **Prompt-level safety instructions** – Quick instructs the model to promote responsible behavior, resist attempts to override safety settings, and avoid exposing internal system information.
- **Access and permission controls** – Identity, data-access, and action-permission controls limit what data and operations are available during an interaction. These limits match the signed-in person's permissions. For more information about identity and permissions, see [Identity and access management in Quick](#) and [Custom permissions](#).
- **Administrator-configured controls** – You can set up blocked words and phrases for chat interactions in the Amazon Quick web experience.

These controls help reduce risk. However, they do not replace your responsibility to configure access, permissions, agents, flows, and data sources for your use case.

Blocked words and phrases

You can block up to 50 words or phrases in chat interactions in the Amazon Quick web experience. This applies to chat agents and flows. Each entry can contain up to 36 characters. Quick checks both user requests and generated responses for these phrases. Quick doesn't block any words or phrases by default.

Add blocked words and phrases

Use the following procedure to add blocked words and phrases for chat agents and flows in the Amazon Quick web experience.

To add blocked words and phrases

1. Sign in to Quick and choose **Manage Quick**.
2. From the navigation pane, choose **Customization**, and then choose **Chat agent customization**.
3. In **Chat agent customization**, under **Guardrails and safety controls**, choose **Add** to add blocked words and phrases. You can add up to 50 words and phrases. Each entry can contain up to 36 characters.

Edit blocked words and phrases

Use the following procedure to add or remove blocked words and phrases for chat agents and flows in the Amazon Quick web experience.

To edit blocked words and phrases

1. Sign in to Quick and choose **Manage Quick**.
2. From the navigation pane, choose **Customization**, and then choose **Chat agent customization**.
3. In **Chat agent customization**, under **Guardrails and safety controls**, choose **Remove** to remove existing entries, or choose **Add** to add entries. You can add up to 50 words and phrases. Each entry can contain up to 36 characters.

Encrypting your Amazon Quick data with AWS Key Management Service customer managed keys

Amazon Quick enables you to encrypt your Amazon Quick data with the keys you have stored in AWS Key Management Service. This provides you with the tools to audit access to data and satisfy regulatory security requirements. If you need to do so, you have the option to immediately lock down access to your data by revoking access to AWS KMS keys. The use of your AWS KMS keys by Amazon Quick is logged in AWS CloudTrail. You can use these events to verify which key protects a resource. These events record AWS KMS key use. Administrators and auditors can use these events to trace when and where your AWS KMS keys were used. For application-level activity in Amazon Quick, see [Incident response, logging, and monitoring in Amazon Quick](#).

To create customer managed KMS keys, you use AWS Key Management Service (AWS KMS) in the same AWS account and AWS Region as the Amazon Quick resource. An Amazon Quick administrator can then use a customer managed KMS key to encrypt your Amazon Quick data and control access.

You can create and manage customer managed KMS keys in the Amazon Quick console or with the Amazon Quick APIs. For more information about creating and managing customer managed KMS keys with the Amazon Quick APIs, see [Key management operations](#).

The following rules apply to using customer managed KMS keys with Amazon Quick resources:

- Amazon Quick doesn't support asymmetric AWS KMS keys.
- You can have multiple customer managed KMS keys and one default customer managed KMS key per AWS account per AWS Region.
- By default, Amazon Quick resources are encrypted with Amazon Quick–native encryption strategies.
- Data currently encrypted by a customer managed KMS key will stay encrypted by the key.
- Accounts configured with a customer managed KMS key do not support chat memory. For more information, see [Memory and response personalization](#).

 Note

If you use AWS Key Management Service with Amazon Quick, you are billed for access and maintenance as described in the [AWS Key Management Service Pricing page](#). In your billing statement, the costs are itemized under AWS KMS and not under Amazon Quick.

Where a customer managed KMS key does not apply, Amazon Quick protects data with AWS owned or service-managed keys.

Database server certificates that are not managed by AWS are the responsibility of the customer and should be signed by a trusted CA. For more information, see [Network and database configuration requirements](#).

Use the following topics to learn more about using customer managed KMS keys with Amazon Quick. To learn more about data encryption in Amazon Quick, see [Data protection in Amazon Quick](#).

Topics

- [Customer managed KMS key scope](#)
- [Amazon Q data key](#)
- [Add a customer managed KMS key to your account](#)
- [Verify the key used by Amazon Quick](#)
- [Changing the default customer managed KMS key](#)
- [Removing the default customer managed KMS key](#)
- [Auditing customer managed KMS key usage in CloudTrail](#)
- [Revoking access to a customer managed KMS key](#)
- [Recovering encrypted Amazon Quick data](#)

Customer managed KMS key scope

Amazon Quick uses resource-specific encryption at rest. Support for a customer managed KMS key, and the way that key is applied, differs by resource type. If a resource type is not listed in the following table, do not assume that a customer managed KMS key applies to it.

Amazon Quick applies customer managed KMS keys in two ways:

- **Account default key.** Participating resource types select the account default key when data is written or when a resource-level key association is created. Changing the default affects future writes or resources that resolve it, but does not automatically re-encrypt existing content. The exact lifecycle is resource-specific.
- **Amazon Q data key.** Amazon Q data uses a single account-level key that is set the first time Amazon Q data is created in your account and **cannot be changed afterward**. To use a customer managed KMS key for Amazon Q data, register it as your account default key *before* you create any Amazon Q data.

Customer managed KMS key support requires an Amazon Quick Enterprise subscription.

A customer managed KMS key must be a symmetric AWS KMS key in the same AWS account and AWS Region as the Amazon Quick resource. You can configure one default customer managed KMS key for each Amazon Quick account and Region.

Your key policy must allow the Amazon Quick service principals that use the key. Your key policy must permit both `quicksight.amazonaws.com` and `qbusiness.amazonaws.com`. This applies

when the same key serves as both your default customer managed KMS key and the Amazon Q data key. Omitting `qbusiness.amazonaws.com` can cause knowledge base sync operations to fail.

Amazon Quick uses AWS KMS grants to encrypt and decrypt your data with a customer managed KMS key, and creates these grants on your behalf. If a key policy or service control policy restricts `kms:CreateGrant` for the Amazon Quick service principals, encryption operations can fail. Retiring a grant removes Amazon Quick access to the key for that resource.

Enabling automatic key rotation in AWS KMS rotates the key material without changing the key ARN, so it does not change how Amazon Quick uses the key and requires no action in Amazon Quick. Rotating to a *different* AWS KMS key is a default-key change; see the following scope table for how each resource type behaves. The Amazon Q data key cannot be moved to a different key.

To see both values for your account, call `DescribeKeyRegistration`. The response returns your registered keys in `KeyRegistration` and your Amazon Q data key in `QDataKey`. The `QDataKey` value is read-only.

Key registration support follows Amazon Quick Region availability. Individual Amazon Quick features and their customer managed KMS key integrations might have narrower Region or account availability; see the documentation for that feature and use `DescribeKeyRegistration` to inspect your account.

For more information about supported AWS Regions, see [AWS Regions, websites, IP address ranges, and endpoints](#).

Amazon Quick resource type	Key used	When it applies	Existing-resource behavior
SPICE dataset	Account default key	A new dataset uses the default key at creation.	To move an existing dataset to a new default key, perform a full refresh.
Report artifact from a dashboard snapshot, scheduled report, export, or dashboard	Account default key	A newly generated artifact uses the current default key.	Existing artifacts remain encrypted with their existing key and are not re-encrypted.

Amazon Quick resource type	Key used	When it applies	Existing-resource behavior
Data-source credentials, OAuth tokens, and OAuth client applications	Account default key	Encrypted at creation with the current default key.	Re-encrypted with the current default key whenever the credential is updated or an OAuth token is refreshed. If no default key is registered, Amazon Quick uses a service-managed key.
Flows – definitions, publication data, run/session state, history, and uploaded-file names	Account default key	Participating Flow content uses the current default key when written.	Existing encrypted content remains tied to the key recorded when it was written; changing the default does not automatically re-encrypt it.
Automate – participating group and trigger content; validated inbound-email content, attachments, and email sender, subject, and attachment-name audit fields	Account default key	Applies only to these participating Automate stores where the integration is available.	Existing records retain the key used for encryption. Other Automate-owned persistence, including workflow/deployment/run data, execution artifacts, and credentials, uses service-controlled keys rather than the account default key.

Amazon Quick resource type	Key used	When it applies	Existing-resource behavior
Research – artifacts , documents, reports/exports, and encrypted record fields	Account default key associated at Research creation	A new Research resource associates the current default key when available ; otherwise Amazon Quick uses service encryption. To use a customer managed KMS key for a Research resource, register it as your default key before creating the resource – a Research resource created without a key association cannot be moved to one later.	Existing Research resources remain associated with their creation-time key when the account default changes. All content in the resource, including new writes, continues to use the creation-time key.
Amazon Q data – knowledge bases, spaces and space uploads, conversations, agents, actions, extensions (including extension access configuration), topics, and Quick Index	Amazon Q data key	Set the first time Amazon Q data is created in the account, from the default key registered at that moment.	Cannot be changed. Not affected when you change or remove the default key.

Amazon Q data key

Amazon Quick protects Amazon Q data with a single account-level key called the Amazon Q data key. Amazon Quick sets this key the first time Amazon Q data is created in your account:

- If a default customer managed KMS key is registered at that moment, the Amazon Q data key is set to that key and `QDataKey.QDataKeyType` is CMK.

- If no default key is registered, Amazon Quick uses an AWS owned key and `QDataKey.QDataKeyType` is `AWS_OWNED`.

The Amazon Q data key cannot be changed after it is set. Changing or removing your default customer managed KMS key does not change it, and does not return an account to an AWS owned key. To use a customer managed KMS key for Amazon Q data, register that key as your account default key before you create any Amazon Q data.

To check the current value, call `DescribeKeyRegistration` and read the `QDataKey` field.

Important

If the key referenced by your Amazon Q data key becomes unavailable, Amazon Q data becomes inaccessible and creating new Amazon Q resources fails. A key can become unavailable if it is disabled, if access is denied, or if it is scheduled for deletion. Because the Amazon Q data key cannot be changed, removing the key registration does not resolve this. Registering a different customer managed KMS key also does not resolve this. The only way to restore Amazon Q functionality is to restore the original key: re-enable it, restore access to it, or cancel its pending deletion. If the key has been permanently deleted, Amazon Q data encrypted with it cannot be recovered. To restore Amazon Q functionality, you must unsubscribe from Amazon Quick and subscribe again. This permanently deletes all of your Amazon Quick data. Before you disable, deny access to, or schedule deletion of any AWS KMS key, call `DescribeKeyRegistration` and confirm the key is not your Amazon Q data key.

Add a customer managed KMS key to your account

Before you begin, make sure that you have an IAM role that grants the admin user access to the Amazon Quick admin key management console. For more information on the required permissions, see [IAM identity-based policies for Amazon Quick: using the admin key management console](#).

You can add keys that already exist in AWS KMS to your Amazon Quick account, so that you can encrypt your Amazon Quick data.

To learn more about how you can create a key to use in Amazon Quick, see the [AWS Key Management Service Developer Guide](#).

To add a new customer managed KMS key to your Amazon Quick account.

1. On the Amazon Quick start page, choose **Manage Amazon Quick**, and then choose **KMS keys**.
2. On the **KMS keys** page, choose **Manage**. The **KMS keys** dashboard opens.
3. On the **KMS Keys** dashboard, choose **Select key**.
4. On the **Select key** pop-up box, choose **Key** to open the list. Then, select the key that you want to add.

If your key isn't in the list, you can manually enter the key's ARN.

5. (Optional) Select the **Use as default encryption key for all new data in the current region of this Amazon Quick account** to set the selected key as your default key. A badge appears next to the default key to indicate its status.

When you choose a default key, new data for participating resource types is encrypted with that key. Support and lifecycle differ by resource type; see [Customer managed KMS key scope](#).

6. (Optional) Add more keys by repeating the previous steps in this procedure. While you can add as many keys as you want, you can only have one default key at one time.

Verify the key used by Amazon Quick

When a key is used, an audit log is created in AWS CloudTrail. You can use the log to track the key's usage. If you need to know which key the Amazon Quick data is encrypted by, you can find this information in CloudTrail.

To learn more about which data can be managed with the key, see [Encrypting your Amazon Quick data with AWS Key Management Service customer managed keys](#).

Verify the customer managed KMS key that's currently used by a SPICE dataset

1. Navigate to your CloudTrail log. For more information, see [Logging Amazon Quick operations with AWS CloudTrail](#).
2. Locate the most recent grant events for the SPICE dataset, using the following search arguments:
 - The event name (eventName) contains Grant.
 - The request parameters requestParameters contain the Amazon Quick ARN for the dataset.

```
{
  "eventVersion": "1.08",
  "userIdentity": {
    "type": "AWSService",
    "invokedBy": "quicksight.amazonaws.com"
  },
  "eventTime": "2022-10-26T00:11:08Z",
  "eventSource": "kms.amazonaws.com",
  "eventName": "CreateGrant",
  "awsRegion": "us-west-2",
  "sourceIPAddress": "quicksight.amazonaws.com",
  "userAgent": "quicksight.amazonaws.com",
  "requestParameters": {
    "constraints": {
      "encryptionContextSubset": {
        "aws:quicksight:arn": "arn:aws:quicksight:us-
west-2:111122223333:dataset/12345678-1234-1234-1234-123456789012"
      }
    },
    "retiringPrincipal": "quicksight.amazonaws.com",
    "keyId": "arn:aws:kms:us-
west-2:111122223333:key/87654321-4321-4321-4321-210987654321",
    "granteePrincipal": "quicksight.amazonaws.com",
    "operations": [
      "Encrypt",
      "Decrypt",
      "DescribeKey",
      "GenerateDataKey"
    ]
  },
  ....
}
```

3. Depending on the event type, one of the following applies:

CreateGrant – You can find the most recently used customer managed KMS key in the key ID (keyID) for the last CreateGrant event for the SPICE dataset.

RetireGrant – If latest CloudTrail event of the SPICE datasets is RetireGrant, there is no key ID and the resource is no longer customer managed KMS key encrypted.

Verify the customer managed KMS key that's currently used when generating report artifacts

1. Navigate to your CloudTrail log. For more information, see [Logging Amazon Quick operations with AWS CloudTrail](#).
2. Locate the most recent GenerateDataKey events for the report execution, using the following search arguments:
 - The event name (eventName) contains GenerateDataKey or Decrypt.
 - The request parameters (requestParameters) contain the Amazon Quick ARN for the analysis or dashboard the report was generated for.

```
{
  "eventVersion": "1.11",
  "userIdentity": {
    "type": "AWSService",
    "invokedBy": "quicksight.amazonaws.com"
  },
  "eventTime": "2025-07-23T23:33:46Z",
  "eventSource": "kms.amazonaws.com",
  "eventName": "GenerateDataKey",
  "awsRegion": "us-west-2",
  "sourceIPAddress": "quicksight.amazonaws.com",
  "userAgent": "quicksight.amazonaws.com",
  "requestParameters": {
    "keyId": "arn:aws:kms:us-west-2:111122223333:key/87654321-4321-4321-4321-210987654321",
    "keySpec": "AES_256",
    "encryptionContext": {
      "aws:quicksight:arn": "arn:aws:quicksight:us-west-2:111122223333:dashboard/1ca456fe-eb34-4250-805c-b1b9350bd164",
      "aws:s3:arn": "arn:aws:s3:::amazon-quick-owned-bucket"
    }
  },
  ...
}
```

3. `aws:s3:arn` is the Amazon Quick owned S3 bucket where your report artifacts are stored.
4. If you no longer see GenerateDataKey, then new report executions are no longer customer managed KMS key encrypted. Existing report artifacts will remain encrypted.

Verify the customer managed KMS key that's currently used for data-source credentials

1. Navigate to your CloudTrail log. For more information, see [Logging Amazon Quick operations with AWS CloudTrail](#).
2. Locate the most recent Encrypt events for the data source, using the following search arguments:
 - The event name (eventName) is Encrypt.
 - The request parameters (requestParameters) contain the encryption context aws:quicksight:arn with the Amazon Quick ARN for the data source.

```
{
  "eventVersion": "1.11",
  "userIdentity": {
    "type": "AWSService",
    "invokedBy": "quicksight.amazonaws.com"
  },
  "eventTime": "2025-07-23T23:33:46Z",
  "eventSource": "kms.amazonaws.com",
  "eventName": "Encrypt",
  "awsRegion": "us-west-2",
  "sourceIPAddress": "quicksight.amazonaws.com",
  "userAgent": "quicksight.amazonaws.com",
  "requestParameters": {
    "keyId": "arn:aws:kms:us-west-2:111122223333:key/87654321-4321-4321-4321-210987654321",
    "encryptionContext": {
      "aws:quicksight:arn": "arn:aws:quicksight:us-west-2:111122223333:datasource/example-data-source-id"
    }
  },
  ...
}
```

3. The keyId on the most recent Encrypt event is the key currently protecting the credential. Credentials are re-encrypted with the current default key whenever the credential is updated or an OAuth token is refreshed, so the most recent event reflects the key in use.

Changing the default customer managed KMS key

You can change the default key to another key that already exists in the **KMS keys** dashboard. When you change the default key, new writes and new resources for participating resource types use the new key. Existing content remains encrypted with the key that was in effect when it was written; how and whether it moves to the new key differs by resource type – for example, a SPICE dataset moves on its next full refresh, data-source credentials move on their next update or token refresh, and report artifacts are never re-encrypted. The Amazon Q data key is not affected. For more information, see [Customer managed KMS key scope](#).

To learn more about which data can be managed with the key, see [Encrypting your Amazon Quick data with AWS Key Management Service customer managed keys](#).

To change the default key to an existing key

1. On the Amazon Quick start page, choose **Manage Amazon Quick**, and then choose **KMS keys**.
2. Choose **MANAGE** to open the **KMS keys** dashboard.
3. Navigate to the key that you want to set as your new default. Choose **Actions** (three dots) on the row of the key that you want to open the key's menu.
4. Choose **Set as default**, and then choose **Set**.

The selected key is now your default key.

Removing the default customer managed KMS key

Removing your default customer managed KMS key ends the use of that key for **new resources or writes that resolve the account default key**. Removing the key does not delete the AWS KMS key, change the key that protects existing resources, change your Amazon Q data key, or remove encryption at rest.

Where no customer managed KMS key applies, Amazon Quick continues to protect data with its own encryption. Existing resources remain encrypted with the key they were created with.

Important

Before you remove your default customer managed KMS key, call `DescribeKeyRegistration` and confirm that the key is not your Amazon Q data key. If the key is your Amazon Q data key, removing it from the default registration does not

change the Amazon Q data key, but future Amazon Q resource behavior depends on continued access to that key. For more information, see [Amazon Q data key](#).

To remove the default customer managed KMS key

1. On the Amazon Quick start page, choose **Manage Amazon Quick**, and then choose **KMS keys**.
2. On the **KMS keys** page, choose **Manage** to open the **KMS keys** dashboard.
3. Choose **Actions** (three dots) on the row of the default key, and then choose **Delete**.
4. In the pop-up box that appears, choose **Remove**.

Auditing customer managed KMS key usage in CloudTrail

You can audit your account's customer managed KMS key usage in AWS CloudTrail. To audit your key usage, log in to your AWS account, open CloudTrail, and choose **Event history**.

Note

Amazon Q data key use appears in CloudTrail as a single account-level key rather than as per-resource grants. `DescribeKeyRegistration` is the authoritative source for the current Amazon Q data key value.

Revoking access to a customer managed KMS key

You can revoke access to your customer managed KMS keys. When you revoke access to a key that is used to encrypt your Amazon Quick data, access to it is denied, and operations that depend on the key can fail until you undo the revoke. The following methods are examples of how you can revoke access:

- Turn off the key in AWS KMS.
- Add a Deny policy to your Amazon Quick AWS KMS policy in IAM.

Important

Before you revoke access to a AWS KMS key, call `DescribeKeyRegistration` and confirm the key is not your Amazon Q data key. If you revoke access to your Amazon Q

data key, Amazon Q data becomes inaccessible and creating new Amazon Q resources fails. This is not reversible by registering a different customer managed KMS key. For more information, see [Amazon Q data key](#).

To learn more about which data can be managed with the key, see [Encrypting your Amazon Quick data with AWS Key Management Service customer managed keys](#).

Use the following procedure to revoke access to your customer managed KMS keys in AWS KMS.

To turn off a customer managed KMS key in AWS Key Management Service

1. Log in to your AWS account, open AWS KMS, and choose **Customer managed keys**.
2. Select the key that you want to turn off.
3. Open the **Key actions** menu and choose **Disable**.

The following is a revocation example that uses a Deny policy in AWS Identity and Access Management (IAM) to prevent further use of the customer managed KMS key. Use "Service": "quicksight.amazonaws.com" as the principal and the ARN of the key as the resource. Deny the following actions: "kms:Encrypt", "kms:Decrypt", "kms:ReEncrypt*", "kms:GenerateDataKey*", "kms:DescribeKey".

Important

After you revoke access by using any method, it can take up to 15 minutes for the data to become inaccessible.

Recovering encrypted Amazon Quick data

The following procedure applies when access to a AWS KMS key has been revoked (for example, by disabling the key or adding a Deny policy), but the key itself still exists in AWS KMS.

Important

If the AWS KMS key has been permanently deleted, data encrypted with it cannot be recovered. AWS KMS enforces a waiting period before key deletion; if the key is still pending deletion, you can cancel the scheduled deletion to restore it. If the permanently

deleted key is your Amazon Q data key, Amazon Q functionality can only be restored by unsubscribing from Amazon Quick and subscribing again, which permanently deletes all of your Amazon Quick data.

To recover Amazon Quick data while its access is revoked

1. Restore access to the customer managed KMS key. Usually, this is enough to recover the Amazon Quick data.
2. Test the Amazon Quick data to check if you can see it.
3. (Optional) If a SPICE dataset is not fully recovered even after you restore access to the customer managed KMS key, perform a full refresh on the dataset. Full refresh applies only to SPICE datasets; it is not applicable to other resource types such as conversations, spaces, agents, or actions.

To learn more about which data can be managed with the key, see [Encrypting your Amazon Quick data with AWS Key Management Service customer managed keys](#).

Permissions

Permission management in Quick controls what actions users can perform and which resources they can access within your Amazon Quick environment. This section covers essential permission-related tasks including configuring access to AWS data sources, setting default access levels for new users, managing IAM policy assignments, creating custom permission profiles tailored to your organization's needs, configuring connectors for third-party integrations, managing extension application access, and integrating with SageMaker AI for machine learning capabilities. Proper permission management ensures users have appropriate access to perform their roles while maintaining security and compliance requirements.

Topics

- [Accessing AWS resources](#)
- [Default access policy](#)
- [IAM policy assignments](#)
- [Extension access](#)
- [Custom permissions](#)

- [Limit profiles](#)

Accessing AWS resources

Applies to: Enterprise Edition and Standard Edition

Intended audience: System administrators and Amazon Quick administrators

You can control the AWS resources that Amazon Quick can access and scope down access to these resources at a more granular level. In Enterprise edition, you can also set up general access defaults for everyone in your account, and you can set up specific access for individual users and groups.

These access configurations are essential for Amazon Quick Sight data source connectivity, enabling secure connections to AWS services like Amazon S3, Amazon RDS, Amazon Redshift, and Athena for data analysis and visualization. Proper resource access setup ensures that Amazon Quick Sight can retrieve and process data from your AWS data sources while maintaining appropriate security boundaries.

Use the following sections to help you configure your AWS resources to work with Quick.

Before you begin, make sure that you have the correct permissions; your system administrator can give you these. To do so, your system administrator creates a policy that enables you to use certain IAM actions. Your system administrator then associates that policy with your user or group in IAM. The required actions are the following:

- **quicksight:AccountConfigurations** – To enable setting default access to AWS resources
- **quicksight:ScopeDownPolicy** – Scoping policies for permissions to AWS resources
- You can also bring your own IAM roles into Amazon Quick. For more information, see [Passing IAM roles to Amazon Quick](#).

To enable or disable the AWS services that Amazon Quick can access

1. Sign in to Amazon Quick at <https://quicksight.aws.amazon.com/>.
2. At the upper right, choose your user name, and then choose **Manage Quick**.
3. Choose **Security & permissions**.

4. Under **QuickSight access to AWS services**, choose **Add or remove**.

A screen appears where you can enable all available AWS services.

 Note

If you see a permissions error, and you're an authorized Amazon Quick administrator, contact your system administrator for assistance.

5. Select the check boxes for the services that you want to allow. Clear check boxes for services that you don't want to allow.

If you have already enabled an AWS service, the check box for that service is already selected. If Amazon Quick can't access a particular AWS service, its check box is not selected.

In some cases, you might see a message like the following.

This policy used by Amazon Quick for AWS resource access was modified outside of Amazon Quick, so you can no longer edit this policy to provide AWS resource permission to Amazon Quick. To edit this policy permissions, go to the IAM console and delete this policy permission with policy arn - `arn:aws:iam::111122223333:policy/service-role/AWSQuickSightS3Policy`.

This type of message means that one of the IAM policies that Amazon Quick uses was manually altered. To fix this, the system administrator needs to delete the IAM policy listed in the error message and reload the **Security & permissions** screen before you try again.

6. Choose **Update** to confirm, or **Cancel** to return to the previous screen.

Topics

- [Setting granular access to AWS services through IAM](#)
- [Using AWS Secrets Manager secrets instead of database credentials in Quick](#)

Setting granular access to AWS services through IAM

Applies to: Enterprise Edition

Intended audience: System administrators and Amazon Quick administrators

In Enterprise edition, Amazon Quick provides a way for you to set up detailed access to resources in AWS services. Like every other AWS service, Quick uses IAM policies to control access for users and groups.

Before you begin, ask an administrator to set up the necessary IAM policies ahead of time. If these are set up, you can select them as part of the procedure in this section. For information about creating IAM policies to use with Quick, see [Identity and access management in Quick](#).

To assign an IAM policy to a user or group

1. Sign in to Quick at <https://quicksight.aws.amazon.com/>.
2. At upper left, choose your user name, and then choose **Manage QuickSight**.
3. Choose **Security & permissions**.
4. Under **Resource access for individual users and groups**, choose **IAM policy assignments**.

The remaining steps at this point involve choosing an IAM policy to assign to the user or group. You can assign multiple IAM policies to one Amazon Quick user or group. To determine permissions, Amazon Quick performs a union and an intersection with the AWS account-level policies.

If you already have active IAM policy assignments, they are listed on this page. You can search for existing assignments by using the search box. If you have drafts that aren't active yet, they are listed under **Assignment drafts**.

5. Choose one of the following:
 - To create an IAM policy assignment, choose **Add new assignment**.
 - To edit an existing assignment, choose the **Edit assignment** icon for that assignment.
 - To enable or disable a policy, select the check box for that policy, and then choose **Enable** or **Disable**. You can select multiple policy assignments at a time.
 - To delete an existing assignment, choose the **Remove assignment** icon near the name of the assignment. To confirm your choice, choose **Delete** on the confirmation screen. Or choose **Back** to cancel deletion.

If you are creating or editing an assignment, continue to the next step. Otherwise, skip to the end of this procedure.

6. On the next screen, you perform the policy assignment process, which is divided into steps. As you work through the steps, you can go forward or backward to make changes. When you exit the screen, your changes from all of the steps are saved.
 - a. **Step 1: Name assignment** – If this is a new assignment, enter a name for the assignment, and then choose **Next** to continue. If you want to change the name, choose **Step 1** at left.
 - b. **Step 2: Select an IAM policy** – Choose an IAM policy that you want to use. From this screen, you can interact with the policies as follows:
 - Choose a policy that you want to use.
 - Search for a policy name.
 - Filter the list to see all IAM policies, AWS-managed policies, or customer-managed policies.
 - View a policy, by choosing **View policy**.

To choose a policy, choose the button beside it, and then choose **Next** to continue.

- c. **Step 3: Assign users and groups** – Choose specific users or groups. Or choose to use the selected IAM policy for all users and groups.

Choose one of the following.

- For **Assign to all users and groups**, select the check box to assign the IAM policy to all Amazon Quick users and groups. Choosing this option assigns the policy to all current and future users and groups.
- Choose the users and groups you want to assign to this IAM policy. You can search for them by name, email address, or group name.

When you are finished selecting users and groups, choose **Next** to continue.

- d. **Step 4: Review and enable changes** – Save your changes.

Choose one of the following.

- To edit any of your choices, choose that step to edit it.

- To save this policy assignment as a draft, choose **Save as draft**. You can enable the draft later.
- To immediately enable this policy, choose **Save and enable**. This option overwrites any existing policy assignment with the same name.

Using AWS Secrets Manager secrets instead of database credentials in Quick

Intended audience: Amazon Quick Administrators and Amazon Quick developers

AWS Secrets Manager is a secret storage service that you can use to protect database credentials, API keys, and other secret information. Using a key helps you ensure that the secret can't be compromised by someone examining your code, because the secret isn't stored in the code. For an overview, see the [AWS Secrets Manager User Guide](#).

Quick administrators can grant Amazon Quick read-only access to secrets they create in Secrets Manager. These secrets can be used in place of database credentials when creating and editing data sources using the Quick API.

Quick supports using secrets with data source types that support credential pair authentication. Jira and ServiceNow are not currently supported.

Note

If you use AWS Secrets Manager with Quick, you are billed for access and maintenance as described in the [AWS Secrets Manager Pricing page](#). In your billing statement, the costs are itemized under Secrets Manager and not under Amazon Quick.

Use the procedures described in the following sections to integrate Secrets Manager with Amazon Quick.

Topics

- [Granting Amazon Quick access to Secrets Manager and selected secrets](#)
- [Creating or updating a data source with secret credentials using the Amazon Quick API](#)
- [What's in the secret](#)

- [Modify a secret](#)

Granting Amazon Quick access to Secrets Manager and selected secrets

If you're an administrator and you have secrets in Secrets Manager, you can grant Amazon Quick read-only access to selected secrets.

To grant Amazon Quick access to Secrets Manager and selected secrets

1. In Amazon Quick, choose your user icon on the upper right, and then choose **Manage Quick**.
2. Choose **Security & permissions** on the left.
3. Choose **Manage** in **Amazon Quick access to AWS resources**.
4. In **Allow access and autodiscovery for these resources**, choose **AWS Secrets Manager, Select secrets**.

The **AWS Secrets Manager secrets** page opens.

5. Select the secrets that you want to grant Amazon Quick read-only access to.

Secrets in your Amazon Quick sign-up Region are shown automatically. To select secrets outside your home Region, choose **Secrets in Other AWS Regions**, and then enter the Amazon Resource Names (ARNs) for those secrets.

6. When you're done, choose **Finish**.

Amazon Quick creates an IAM role called `aws-quicksight-secretsmanager-role-v0` in your account. It grants users in the account read-only access to the specified secrets and looks similar to the following:

When Amazon Quick users create analyses from or view dashboards that use a data source with secrets, Amazon Quick assumes this Secrets Manager IAM role. For more information about secret permissions policies, see [Authentication and access control for AWS Secrets Manager](#) in the *AWS Secrets Manager User Guide*.

The specified secret in the Amazon Quick IAM role may have an additional resource policy that denies access. For more information, see [Attach a permissions policy to a secret](#) in the *AWS Secrets Manager User Guide*.

If you're using an AWS managed AWS KMS key to encrypt your secret, Amazon Quick doesn't require any additional permissions setup in Secrets Manager.

If you're using a customer managed key to encrypt your secret, ensure that the Amazon Quick IAM role, `aws-quicksight-secretsmanager-role-v0` has `kms:Decrypt` permissions. For more information, see [Permissions for the KMS key](#) in the *AWS Secrets Manager User Guide*.

For more information about the types of keys used in AWS Key Management Service, see [Customer keys and AWS keys](#) in the *AWS Key Management Service guide*.

Creating or updating a data source with secret credentials using the Amazon Quick API

After the Amazon Quick administrator has granted Amazon Quick read-only access to Secrets Manager, you can create and update data sources in the API using a secret the administrator selected as credentials.

Following is an example API call to create a data source in Amazon Quick. This example uses the `create-data-source` API operation. You can also use the `update-data-source` operation. For more information, see [CreateDataSource](#) and [UpdateDataSource](#) in the *Amazon Quick API Reference*.

The user specified in the permissions in the following API call example can delete, view, and edit data sources for the specified MySQL data source in Amazon Quick. They can also view and update the data source permissions. Instead of a Amazon Quick username and password, a secret ARN is used as credentials for the data source.

```
aws quicksight create-data-source
  --aws-account-id AWSACCOUNTID \
  --data-source-id DATASOURCEID \
  --name NAME \
  --type MYSQL \
  --permissions '[{"Principal": "arn:aws:quicksight:region:accountID:user/namespace/username", "Actions": ["quicksight:DeleteDataSource",
    "quicksight:DescribeDataSource", "quicksight:DescribeDataSourcePermissions",
    "quicksight:PassDataSource", "quicksight:UpdateDataSource",
    "quicksight:UpdateDataSourcePermissions"]}]' \
  --data-source-parameters='{"MySQLParameters":{"Database": "database",
    "Host": "hostURL", "Port": "port"}}' \
  --
credentials='{"SecretArn": "arn:aws:secretsmanager:region:accountID:secret:secretname"}' \
  --region us-west-2
```

In this call, Amazon Quick authorizes `secretsmanager:GetSecretValue` access to the secret based on the API caller's IAM policy, not the IAM service role's policy. The IAM service role acts on the account level and is used when an analysis or dashboard is viewed by a user. It cannot be used to authorize secret access when a user creates or updates the data source.

When they edit a data source in the Amazon Quick UI, users can view the secret ARN for data sources that use AWS Secrets Manager as the credential type. However, they can't edit the secret, or select a different secret. If they need to make changes, for example to the database server or port, users first need to choose **Credential pair** and enter their Amazon Quick account username and password.

Secrets are automatically removed from a data source when the data source is altered in the UI. To restore the secret to the data source, use the `update-data-source` API operation.

What's in the secret

Amazon Quick requires the following JSON format to access your secret:

```
{
  "username": "username",
  "password": "password"
}
```

The `username` and `password` fields are required for Amazon Quick to access secrets. All other fields are optional and are ignored by Amazon Quick.

The JSON format may vary depending on the type of database. For more information, see [JSON structure of AWS Secrets Manager database credential secrets](#) in the *AWS Secrets Manager User Guide*.

Modify a secret

To modify a secret, you use Secrets Manager. After you make changes to a secret, the updates become available the next time Amazon Quick requests access to the secret.

Default access policy

Applies to: Enterprise Edition

Intended audience: System administrators and Amazon Quick administrators

In the Enterprise edition, you can configure specific permissions for the AWS services that an Amazon Quick user can access. If no such configuration occurs, Quick uses a default set of permissions based on the user's settings. The current behavior is displayed in a blue information box.

To change the default resource access for all users (to use when no other permissions are configured)

1. Sign in to Amazon Quick.
2. At upper right, choose the profile icon, and then choose **Manage Quick**.
3. Under **Permissions**, choose **Default access policy**.
4. Choose one of the following:
 - Allow access to all AWS data and resources to all users and groups.
 - Deny access to all AWS data and resources to all users and groups.
5. Choose **Update**.

IAM policy assignments

IAM policy assignments allow you to control access to AWS resources for users and groups within Amazon Quick. By assigning IAM policies, you can define and manage permissions for various AWS services and resources that Quick may need to access. Assignments can be enabled, disabled, or deleted at any time.

To add a new IAM policy assignment:

1. Open the [Quick console](#).
2. Choose the user icon at the top right, and then choose **Manage Quick**.
3. Under the **Permissions** section, choose **IAM policy assignments**.
4. Choose **Add new assignment**.
5. Enter a name for this assignment and choose **Next**.
6. Select an IAM policy by either searching for it directly or scrolling through the displayed list, and then choose **Next**.

7. Assign users or groups to the IAM policy by either searching by name, email address, or group name, or scrolling through the displayed list and selecting the appropriate users or groups. You can also choose to assign the policy to all users and groups by selecting the checkbox; this will apply the policy to all current and future users and groups.
8. Choose **Next**.
9. Review your selections and choose **Save and enable**.
10. You can enable or disable the assignment at any time by selecting the checkbox next to the assignment name and choosing **Enable** or **Disable**. You can view the status of the policy assignment under the **Status** column.

To edit an IAM policy assignment:

1. Open the [Quick console](#).
2. Choose the user icon at the top right, and then choose **Manage Quick**.
3. Under the **Permissions** section, choose **IAM policy assignments**.
4. Choose the edit (pencil) icon.
5. Review the sections and choose **Edit** to make any changes.
6. Choose **Save and enable** to apply your updates.

To delete an IAM policy assignment:

1. Open the [Quick console](#).
2. Choose the user icon at the top right, and then choose **Manage Quick**.
3. Under the **Permissions** section, choose **IAM policy assignments**.
4. Choose the delete (trash) icon.
5. A confirmation modal will appear to ensure you want to proceed with the deletion. Choose **Delete** to delete the IAM policy assignment.

Extension access

Extension access connects your Amazon Quick account to third party applications like Microsoft Teams, Microsoft Outlook, Microsoft Word, Microsoft Excel, Microsoft PowerPoint, and Slack. Think of it as giving permission for Amazon Quick to work inside the apps your team already uses, helping them accomplish their work, assisted by generative AI, and without context switching.

As an admin, you configure extension access. Then, either an admin or an author must install the extension in the Amazon Quick console to make the extension available to end users. This default extension is available for use to all users on eligible subscriptions after setup. Installed extensions are available under **Connections > Extensions** in the Amazon Quick menu for further configuration and management.

If you're an admin user of Amazon Quick, you configure Amazon Quick access to extensions, and can also deploy extensions for end user use. As an admin, you can also allow authors to edit and deploy extensions after you configure access.

User roles and extension permissions work as follows:

- **IAM Administrators:** Can enable/remove extension access, create extensions, assign owners, and control extension availability for the organization.
- **Authors:** Can create and install extensions using administrator-enabled access, and edit/delete extensions they own.
- **Readers (Reader Pro only):** Can use all enabled extensions but cannot create or modify them.

Here's what you need to know about extension access:

- **What it is:** A secure connection between your Amazon Quick instance and your organization's account in third-party apps (like your Microsoft 365 or Slack workspace).
- **Why it matters:** It gives your Amazon Quick instance permission to work within your organization's specific tenant or workspace in these apps.
- **Who sets it up:** Only administrators with IAM credentials can configure extension access and deploy extensions (or give authors permission to edit and deploy extensions) - it's a privileged operation that must be done before your team can start using extensions. Note that IAM admin users can access the **Extension access** pages from the **Manage Quick** menu, but to configure a extension they may need to re-log in to the system and select **Connections > Extensions** from the Amazon Quick menu.

Topics

- [Browser extension](#)
- [Amazon Quick Microsoft Outlook extension](#)
- [Amazon Quick Microsoft Excel extension](#)
- [Amazon Quick Microsoft Word extension](#)

- [Amazon Quick Microsoft PowerPoint extension](#)
- [Amazon Quick Slack extension](#)
- [Amazon Quick Microsoft Teams extension](#)

Browser extension

The Amazon Quick browser extension is available within Amazon Quick to all eligible users by default and requires no administrative setup for user access. Users can find the browser extension installation link under **Connections > Extensions**. The following procedures are for IT administrators who want to automatically deploy or disable the Amazon Quick browser extension across all browsers in their organization using managed device policies.

Topics

- [Deploy the browser extension](#)
- [Disable the browser extension](#)

Deploy the browser extension

Note

The following deployment guide applies only when you have a managed device fleet with managed browsers.

To deploy the browser extension to all users in your organization, you can define an **ExtensionSettings** policy and distribute this policy to your managed devices.

Topics

- [Install using ExtensionSettings policy](#)
- [Distribute ExtensionSettings policy](#)

Install using ExtensionSettings policy

The **ExtensionSettings** policy is an administrative setting for Chrome, Edge, and Firefox that allows you to manage the installation, permissions, and runtime behavior of specific extensions on managed devices. This policy provides granular control by letting administrators specify custom

configurations, such as setting installation modes like force install and restricting or allowing specific site access and permissions for individual extensions or all extensions.

The format of the **ExtensionSettings** policy depends on the operating system where you want to distribute this policy. Windows, Mac, and Linux are supported.

To set the install mode in the **ExtensionSettings** policy, you need the extension identifier and the store URL:

- **Chrome/Edge**

- ID: innkphffipcmiflfibbeghfnkifiokgo
- URL: <https://clients2.google.com/service/update2/crx>

- **Firefox**

- ID: quick-browser-extension@amazon.com
- URL: <https://addons.mozilla.org/firefox/downloads/latest/amazon-quick/latest.xpi>

Distribute ExtensionSettings policy

Distributing the **ExtensionSettings** policy depends on the device and browser management solution used by your organization. You can use managed solutions like **Chrome Enterprise Core** or **Microsoft Edge** management service, or distribute the **ExtensionSettings** policy using **Group Policy**, which differs per operating system.

Disable the browser extension

Note

The following deployment guide applies only when you have a managed device fleet with managed browsers.

To disable the browser extension for all users in your organization, you can define an **ExtensionSettings** policy and distribute this policy to your managed devices.

Topics

- [Disable using ExtensionSettings policy](#)

- [Distribute ExtensionSettings policy](#)

Disable using ExtensionSettings policy

The **ExtensionSettings** policy is an administrative setting for Chrome, Edge, and Firefox that allows you to manage the installation, permissions, and runtime behavior of specific extensions on managed devices. This policy provides granular control by letting administrators specify custom configurations, such as disabling individual extensions.

The format of the **ExtensionSettings** policy depends on the operating system where you want to distribute this policy. See the documentation for the **ExtensionSettings** policy to determine the format. Windows, Mac, and Linux are supported.

To block the use of the extension in the **ExtensionSettings** policy, you need the extension identifier:

- **Chrome/Edge**

ID: innkphffipcmiflfibbeghfnkifiokgo

- **Firefox**

ID: quick-browser-extension@amazon.com

Distribute ExtensionSettings policy

Distributing the **ExtensionSettings** policy depends on the device and browser management solution used by your organization. You can use managed solutions like **Chrome Enterprise Core** or **Microsoft Edge** management service, or distribute the **ExtensionSettings** policy using **Group Policy**, which differs per operating system.

Amazon Quick Microsoft Outlook extension

The Quick extension for Microsoft Outlook integrates AI-powered assistance directly into your email workflows. Users can leverage Quick within Outlook to streamline their inbox and calendar management, draft contextual emails using their Quick data, and perform external actions without switching applications.

The Outlook extension enables users to:

- Prioritize and organize their inbox.

- Search for relevant email content for a topic using natural language.
- Schedule meetings and manage their calendars.
- Summarize emails and draft contextual responses using Quick data.
- Access their enterprise data and perform actions in external applications.

The Amazon Quick Outlook extension is available within Amazon Quick to all eligible users by default and requires no administrative setup for user access if users can access the Microsoft Outlook app store. Users can search for Amazon Quick in the Microsoft Outlook app store or visit the [Quick for Outlook Microsoft store page](#) to add the extension.

Important

- The Amazon Quick Outlook extension uses generative AI to create and execute code within your Outlook application sandbox to help you perform your Outlook tasks. AI can make mistakes and perform inaccurate actions within your Outlook mailbox.
- Amazon Quick does not use your user data for service improvement or for training its underlying large language models (LLMs).

The following procedures are for IT administrators who want to automatically deploy the Amazon Quick Outlook extension across their organization on behalf of their users.

Topics

- [Prerequisites for deploying the Microsoft Outlook extension to your organization](#)
- [Microsoft Outlook extension permissions](#)
- [Deploying the Microsoft Outlook extension to your organization](#)

Prerequisites for deploying the Microsoft Outlook extension to your organization

Before configuring access to the Amazon Quick extension to Microsoft Outlook, administrators must complete the following steps:

1. Have a Microsoft 365 Business subscription and be a Global Admin or have administrative permissions (specifically `AppCatalog.ReadWrite.All`).
2. Have an Amazon Quick instance.

3. Grant tenant-wide admin consent for Microsoft Graph permissions. This allows the extension to access inbox, calendar, and contact data for all users in your organization. Without these permissions, the extension can only respond to questions about the currently open email thread. For instructions, see [the section called "Grant tenant-wide admin consent"](#).

Microsoft Outlook extension permissions

The Amazon Quick Outlook extension uses the ReadWriteMailbox Office JavaScript API permission level. This is an Office add-in manifest permission, not a Microsoft Graph API scope.

The default app capabilities for the Outlook add-in are:

- Can read and make changes to your document
- Access your profile information such as your name, email address, company name, and preferred language
- Can send data over the Internet

Graph API delegated permissions

The Outlook extension uses Microsoft Graph integration, which requires additional delegated permissions.

Important

These permissions are required for the Amazon Quick Outlook extension to provide its full range of capabilities. Without these permissions granted, the extension can only answer questions about the currently open email thread and draft a reply. It cannot summarize your inbox, manage your calendar, or perform inbox-wide operations such as searching, organizing, or triaging emails.

Note

The Amazon Quick Outlook extension does not send emails or create calendar appointments automatically. It uses Human-in-the-Loop (HITL) confirmation before any sending action, so users always review and approve outgoing messages and calendar changes.

Users are prompted to sign in the first time they interact with the extension. Administrators can pre-approve these permissions by granting tenant-wide admin consent (see [the section called "Grant tenant-wide admin consent"](#)).

The following table lists the delegated Graph API permissions used by the Outlook extension.

Permission	Purpose
<code>Calendars.Read.Shared</code>	Read events in user's and shared/delegate calendars for availability checks
<code>Calendars.ReadWrite</code>	Create, update, and delete calendar events; accept/decline invites; check free/busy times
<code>Contacts.ReadWrite</code>	Resolve recipients by name, look up contact details, and save new contacts
<code>Files.Read</code>	Read user's OneDrive/SharePoint files to summarize or reference attachments
<code>Mail.Read</code>	Read full email content, attachments, and threads for summarization and action items
<code>Mail.ReadBasic</code>	Read email metadata (subject, sender, date) for efficient mailbox listing and search
<code>Mail.ReadWrite</code>	Organize mail: move, delete, mark read/unread, create drafts, and manage attachments
<code>MailboxSettings.ReadWrite</code>	Read and update timezone, working hours, out-of-office, inbox rules, and categories
<code>offline_access</code>	Obtain refresh token for persistent sessions without repeated re-authentication
<code>People.Read</code>	Suggest recipients and resolve names using collaboration and org signals
<code>Tasks.ReadWrite</code>	Create, update, complete, and delete tasks and task lists across To Do/Planner

Permission	Purpose
User.Read	Identify the signed-in user's profile, email, job title, and timezone
User.ReadBasic.All	Read basic profile info (name, email, photo) of other users in the org

Grant tenant-wide admin consent

To enable the full capabilities of the Amazon Quick Outlook extension for all users in your organization, grant tenant-wide admin consent for the Microsoft Graph permissions listed above. This pre-approves the permissions so that individual users are not prompted to approve them when they first use the extension.

You can complete this process using either Microsoft Graph PowerShell or Microsoft Graph Explorer.

Prerequisites

Before granting admin consent, ensure you have a role that includes the following permissions:

- Application.ReadWrite.All
- DelegatedPermissionGrant.ReadWrite.All

Grant admin consent using PowerShell

Follow these steps to grant admin consent using Microsoft Graph PowerShell:

1. Connect to Microsoft Graph with the required permissions:

```
Connect-MgGraph -Scopes  
"Application.ReadWrite.All","DelegatedPermissionGrant.ReadWrite.All"
```

2. Look up the Microsoft Graph service principal ID for your tenant and save the ID:

```
Get-MgServicePrincipal -Filter "displayName eq 'Microsoft Graph'"
```

3. Create a service principal for the Amazon Quick Outlook extension in your tenant and save the returned ID:

```
New-MgServicePrincipal -AppId "a5342f89-ebb1-4b1d-966c-34e8df972aaf"
```

- Grant admin consent for all delegated permissions on behalf of all users in your tenant. Replace *ServicePrincipalId* with the ID from step 3 and *GraphServicePrincipalId* with the ID from step 2:

```
$params = @{
    "ClientId"      = "ServicePrincipalId"
    "ConsentType"   = "AllPrincipals"
    "ResourceId"    = "GraphServicePrincipalId"
    "Scope"         = "Calendars.Read.Shared Calendars.ReadWrite Contacts.ReadWrite
Files.Read Mail.Read Mail.ReadBasic Mail.ReadWrite MailboxSettings.ReadWrite
offline_access People.Read Tasks.ReadWrite User.Read User.ReadBasic.All"
}
New-MgOauth2PermissionGrant -BodyParameter $params
```

Grant admin consent using Graph Explorer

Follow these steps to grant admin consent using Microsoft Graph Explorer:

- Sign in to Microsoft Graph Explorer with a role that has the `Application.ReadWrite.All` and `DelegatedPermissionGrant.ReadWrite.All` permissions.
- Look up the Microsoft Graph service principal ID for your tenant and save the ID from the response:

```
GET https://graph.microsoft.com/v1.0/servicePrincipals?$filter=displayName eq
'Microsoft Graph'&$select=id,displayName
```

- Create a service principal for the Amazon Quick Outlook extension in your tenant and save the returned ID:

```
POST https://graph.microsoft.com/v1.0/servicePrincipals
```

Request body:

```
{
    "appId": "a5342f89-ebb1-4b1d-966c-34e8df972aaf"
```

```
}
```

4. Grant admin consent for all delegated permissions on behalf of all users in your tenant. Replace *ServicePrincipalId* with the ID from step 3 and *GraphServicePrincipalId* with the ID from step 2:

```
POST https://graph.microsoft.com/v1.0/oauth2PermissionGrants
```

Request body:

```
{
  "clientId": "ServicePrincipalId",
  "consentType": "AllPrincipals",
  "resourceId": "GraphServicePrincipalId",
  "scope": "Calendars.Read.Shared Calendars.ReadWrite Contacts.ReadWrite
Files.Read Mail.Read Mail.ReadBasic Mail.ReadWrite MailboxSettings.ReadWrite
offline_access People.Read Tasks.ReadWrite User.Read User.ReadBasic.All"
}
```

After granting admin consent, users in your organization will not be prompted to approve permissions individually when they first use the Amazon Quick Outlook extension. The extension will have immediate access to its full capabilities.

Where to review permissions

To review permissions for the Outlook extension, go to Microsoft Entra > **Enterprise Applications** > find the app called "Amazon Quick Outlook Agent" > **Permissions**.

Deploying the Microsoft Outlook extension to your organization

After completing the prerequisites and granting admin consent for Microsoft Graph permissions, follow these steps to deploy the extension to your users:

1. Sign in to the M365 admin center.
2. Select **Settings** > **Integrated apps** in the left navigation menu.
3. Choose **Get apps**.
4. Search for "Amazon Quick".
5. Locate the tile for Amazon Quick in Outlook and choose **Get it now**.

6. Confirm that you want to add the app.
7. Under **Assign users**, choose **Entire organization** or **Specific users/groups** depending on your needs.
8. After selecting the users, review the app's requested permissions and capabilities and choose **Next**.
9. Choose **Finish Deployment**.

Amazon Quick Microsoft Excel extension

The Quick extension for Microsoft Excel integrates AI-powered assistance directly into your spreadsheet workflows. Users can leverage Quick within Excel to perform analysis, import data and build workbooks, and streamline spreadsheet operations like creation of tables, charts, and formatting.

The Excel extension enables users to:

- Perform in-depth analysis of their workbooks and generate insights.
- Import, clean up, and format data from Quick's enterprise knowledge and apps.
- Generate tables and charts.
- Automate spreadsheet operations such as formatting and applying filters and transformations.
- Apply Excel functions and build and decode complex workbooks.

The Quick Excel extension is available within Amazon Quick to all eligible users by default and requires no administrative setup for user access if users can access the Microsoft Excel app store. Users can search for Amazon Quick in the Microsoft Excel app store or visit the [Quick for Excel Microsoft store page](#) to add the extension.

Important

- The Amazon Quick Excel extension uses generative AI to create and execute code within your Excel application sandbox to help you perform tasks within Excel. AI can make mistakes and perform inaccurate actions within your Excel workbook.
- Amazon Quick does not use your user data for service improvement or for training its underlying large language models (LLMs).

The following procedures are for IT administrators who want to automatically deploy the Amazon Quick Excel extension across their organization on behalf of their users.

Topics

- [Prerequisites for deploying the Microsoft Excel extension to your organization](#)
- [Deploying the Microsoft Excel extension to your organization](#)
- [Microsoft Excel extension permissions](#)

Prerequisites for deploying the Microsoft Excel extension to your organization

Before configuring access to the Amazon Quick extension to Microsoft Excel, administrators must complete the following steps:

1. Have a Microsoft 365 Business subscription and be a Global Admin or have administrative permissions (specifically `AppCatalog.ReadWrite.All`).
2. Have an Amazon Quick instance.

Deploying the Microsoft Excel extension to your organization

Follow these steps to deploy the extension to your users:

1. Login to M365 admin center.
2. Select **Settings** > **Integrated apps** in the left navigation menu.
3. Click on **Get apps**.
4. Search for "Amazon Quick".
5. Locate the tile for the Amazon Quick in Excel and click on **Get it now**.
6. Confirm that you want to add the app.
7. Under **Assign users** you can choose **Entire organization** or **Specific users/groups** depending on your needs.
8. After selecting the users, review the app's requested permissions and capabilities and click **Next**.
9. Click **Finish Deployment**.

Microsoft Excel extension permissions

The Amazon Quick Excel extension uses the ReadWriteDocument Office JavaScript API permission level. This is an Office add-in manifest permission, not a Microsoft Graph API scope. No Graph API permissions are needed during deployment.

The default app capabilities for the Excel add-in are:

- Can read and make changes to your document
- Access your profile information such as your name, email address, company name, and preferred language
- Can send data over the Internet

Administrators should review the permissions listed during the consent step of deployment through the M365 admin center integrated apps portal.

Amazon Quick Microsoft Word extension

The Quick extension for Microsoft Word integrates AI-powered assistance directly in your Word document. Users can leverage Quick within Word to draft content, make formatted edits, and perform reviews with track changes and comments enabled.

The Word extension enables users to:

- Generate formatted documents using style guides and templates and Quick's enterprise data and apps.
- Make small- or large-scale edits to your document while maintaining its writing style and structure.
- Format or restructure your documents.
- Redline your documents with track changes and/or comments enabled.

The Quick Word extension is available within Amazon Quick to all eligible users by default and requires no administrative setup for user access if users can access the Microsoft Word app store. Users can search for Amazon Quick in the Microsoft Word app store or visit the [Quick for Word Microsoft store page](#) to add the extension.

Important

- The Amazon Quick Word extension uses generative AI to create and execute code within your Word application sandbox to help you perform tasks within Word. AI can make mistakes and perform inaccurate actions within your Word document.
- Amazon Quick does not use your user data for service improvement or for training its underlying large language models (LLMs).

The following procedures are for IT administrators who want to automatically deploy the Amazon Quick Word extension across their organization on behalf of their users.

Topics

- [Prerequisites for deploying the Microsoft Word extension to your organization](#)
- [Deploying the Microsoft Word extension to your organization](#)
- [Microsoft Word extension permissions](#)

Prerequisites for deploying the Microsoft Word extension to your organization

Before configuring access to the Amazon Quick extension to Microsoft Word, administrators must complete the following steps:

1. Have a Microsoft 365 Business subscription and be a Global Admin or have administrative permissions (specifically `AppCatalog.ReadWrite.All`).
2. Have an Amazon Quick instance.

Deploying the Microsoft Word extension to your organization

Follow these steps to deploy the extension to your users:

1. Login to M365 admin center.
2. Select **Settings** > **Integrated apps** in the left navigation menu.
3. Click on **Get apps**.
4. Search for "Amazon Quick".
5. Locate the tile for the Amazon Quick in Word and click on **Get it now**.

6. Confirm that you want to add the app.
7. Under **Assign users**, you can choose **Entire organization** or **Specific users/groups** depending on your needs.
8. After selecting the users, review the app's requested permissions and capabilities and click **Next**.
9. Click **Finish Deployment**.

Microsoft Word extension permissions

The Amazon Quick Word extension uses the ReadWriteDocument Office JavaScript API permission level. This is an Office add-in manifest permission, not a Microsoft Graph API scope. No Graph API permissions are needed during deployment.

The default app capabilities for the Word add-in are:

- Can read and make changes to your document
- Access your profile information such as your name, email address, company name, and preferred language
- Can send data over the Internet

Administrators should review the permissions listed during the consent step of deployment through the M365 admin center integrated apps portal.

Amazon Quick Microsoft PowerPoint extension

The Amazon Quick extension for PowerPoint integrates AI-powered assistance directly in your slide deck. Users can leverage Quick within PowerPoint to create presentations, add new slides, and make formatted edits to slides using their Quick enterprise data and apps.

The PowerPoint extension enables users to:

- Create visually appealing and formatted templates using their own templates and brand guidelines.
- Create and edit slides and visuals within PowerPoint using natural language and their Quick enterprise data and apps.
- Make decks presentation ready by streamlining operations such as restructuring presentations, applying transitions, and more.

The Quick PowerPoint extension is available within Amazon Quick to all eligible users by default and requires no administrative setup for user access if users can access the Microsoft PowerPoint app store. Users can search for Amazon Quick in the Microsoft PowerPoint app store or visit the [Quick for PowerPoint Microsoft store page](#) to add the extension.

Important

- The Amazon Quick PowerPoint extension uses generative AI to create and execute code within your PowerPoint application sandbox to help you perform tasks within PowerPoint. AI can make mistakes and perform inaccurate actions within your PowerPoint slide deck.
- Amazon Quick does not use your user data for service improvement or for training its underlying large language models (LLMs).

The following procedures are for IT administrators who want to automatically deploy the Amazon Quick PowerPoint extension across their organization on behalf of their users.

Topics

- [Prerequisites for deploying the Microsoft PowerPoint extension to your organization](#)
- [Deploying the Microsoft PowerPoint extension to your organization](#)
- [Microsoft PowerPoint extension permissions](#)

Prerequisites for deploying the Microsoft PowerPoint extension to your organization

Before configuring access to the Amazon Quick extension to Microsoft PowerPoint, administrators must complete the following steps:

1. Have a Microsoft 365 Business subscription and be a Global Admin or have administrative permissions (specifically `AppCatalog.ReadWrite.All`).
2. Have an Amazon Quick instance.

Deploying the Microsoft PowerPoint extension to your organization

Follow these steps to deploy the extension to your users:

1. Login to M365 admin center.

2. Select **Settings** > **Integrated apps** in the left navigation menu.
3. Click on **Get apps**.
4. Search for "Amazon Quick".
5. Locate the tile for the Amazon Quick in PowerPoint and click on **Get it now**.
6. Confirm that you want to add the app.
7. Under **Assign users**, you can choose **Entire organization** or **Specific users/groups** depending on your needs.
8. After selecting the users, review the app's requested permissions and capabilities and click **Next**.
9. Click **Finish Deployment**.

Microsoft PowerPoint extension permissions

The Amazon Quick PowerPoint extension uses the ReadWriteDocument Office JavaScript API permission level. This is an Office add-in manifest permission, not a Microsoft Graph API scope. No Graph API permissions are needed during deployment.

The default app capabilities for the PowerPoint add-in are:

- Can read and make changes to your document
- Access your profile information such as your name, email address, company name, and preferred language
- Can send data over the Internet

Administrators should review the permissions listed during the consent step of deployment through the M365 admin center integrated apps portal.

Amazon Quick Slack extension

The Quick extension for Slack integrates AI-powered assistance directly into your team communication workflows. Users can access Quick knowledge and capabilities within Slack channels and direct messages without switching between applications.

The Slack extension enables users to:

- Add Amazon Quick as a collaborator using **@Amazon Quick** mentions.
- Get conversation summaries and insights.

- Access organizational knowledge bases and documents directly from Slack channels.
- Generate notes and action items from discussions in Slack threads.
- Get help with data analysis and report generation using uploaded files.

Important

- When Amazon Quick is used in public Slack channels, responses are based on the invoking user's permissions. This may include content that other channel members aren't authorized to access. Carefully evaluate using Amazon Quick in public channels to prevent unintended exposure of sensitive information.
- Amazon Quick does not use your user data for service improvement or for training its underlying large language models (LLMs).

Topics

- [Prerequisites for Slack extension](#)
- [Configure Slack extension access](#)

Prerequisites for Slack extension

Before adding the Amazon Quick Slack Extension, administrators must complete the following requirements:

- Have a paid Slack workspace.
- Have admin access to your Slack workspace.
- Get started with Amazon Quick.
- Your Slack workspace ID (must start with 'T' and be alphanumeric). One way to find your Slack workspace ID is by navigating to your Slack workspace and starting a chat with the Slack Developer Tools app running the `/sdt whoami` command. For more information, see [Locate your Slack URL or ID](#) in the Slack help center.

If you configured the authentication to connect to Amazon Quick with IAM Identity Center, complete the following additional steps:

1. Ensure you have an IAM Identity Center instance enabled on your AWS account.

2. If you are connecting an external IAM provider to IAM Identity Center, ensure that every user under your IAM provider configuration has an email associated with them.
3. If you are using Entra ID, set up SCIM identity propagation between the Microsoft Entra ID instance and IAM Identity Center. For detailed steps, see [Configure SAML and SCIM with Microsoft Entra ID and IAM Identity Center](#).

Configure Slack extension access

As an administrator, you must allow your Amazon Quick Slack extension to connect to your Amazon Quick application environment. You can use the Amazon Quick console to manage extension access configurations.

User attribute mapping

When you configure a Slack extension, user identity is mapped by default using the following attributes:

- **Amazon Quick user attribute** - Email address is used to map Amazon Quick users to their corresponding Slack accounts. The system uses the email address to establish the connection between user identities.
- **Slack user attribute** - User Profile Email is used to match against Slack user accounts. This maps to the email address associated with the user's Slack profile.

These default mappings ensure secure and accurate user identification across both platforms without requiring additional configuration.

Topics

- [Add Slack extension access for accounts using IAM Identity Center](#)
- [Add Slack extension access for accounts using other authentication methods](#)
- [Edit Slack extension access](#)
- [Delete Slack extension access](#)

Add Slack extension access for accounts using IAM Identity Center

Configuring extension access with IAM Identity Center requires completing steps specific to your identity provider (Entra ID or Okta) followed by common setup steps in AWS.

Configure IAM Identity Center with Entra ID

Follow these steps only if you are using IAM Identity Center with Entra ID to set up and configure an Azure tenant on your Microsoft Azure portal:

To set up an Azure tenant

1. In the Azure account, create a new app registration.
 - a. Go to **App registrations**.
 - b. In the **App registrations** screen, choose **New registration**. Under the **Supported account types** option, choose **Accounts in this organizational directory only (Personal use only - Single tenant)**. Once finished, choose **Register**.
 - c. Note the client ID. You will need this later.
 - d. Create a client secret for the app registration and keep note of it. You will need this later.
2. Add callback URLs for each Region in which your Slack extension will be installed.
 - a. Navigate to the app registration's **Authentication** tab.
 - b. Choose **Platform Configurations, Add a platform**.
 - c. Choose **Web**.
 - d. Compose a callback URL using the following format, replacing *your-region* with your Amazon Quick instance Region. The Slack extension supports the following Regions: ap-southeast-2, eu-west-1, us-west-2, and us-east-1.

```
qbs-cell1001.dp.appintegrations.your-region.prod.plato.ai.aws.dev/auth/idc-tti/callback
```
 - e. Insert the callback URL as the redirect URI and choose **Configure** when done.

Follow these steps to configure a Trusted Token Issuer on your IAM Identity Center instance in your AWS Console:

To configure a Trusted Token Issuer

1. Go to your AWS account and navigate to your IAM Identity Center instance.
2. Navigate to **Settings, Authentication**.
3. Choose **Create trusted token issuer**.

4. Add the issuer URL, which should follow this template, where *Tenant ID* refers to your Entra tenant ID:

```
login.microsoftonline.com/Tenant ID/v2.0
```

 Note

The issuer URL should be the OIDC discovery endpoint of your identity without the well-known document URI path. If you include the well-known document URI path, this will not work. See Trusted token issuer configuration settings.

5. Choose **Email** as the Identity Provider attribute and IAM Identity Center attribute.
6. Note the trusted token issuer ARN. You will need this in a later step.

After completing these Entra ID-specific steps, proceed to the [Complete AWS Configuration \(all providers\)](#) section below.

Configure IAM Identity Center with Okta

Follow these steps only if you are using IAM Identity Center with Okta to set up and configure your App Integration in the Okta Admin console:

To set up an Okta Application

1. In your Okta account, create a new Okta App Integration.
 - a. In your Okta Admin console, navigate to **Applications > Applications**.
 - b. Click on **Create App Integration**.
 - c. For the Sign-in method, select **OIDC - OpenID Connect**.
 - d. For the Application type, select **Web Application**.
 - e. Click on **Next**.
 - f. Provide an App integration name.
 - g. Under **Grant type > Core grants**, ensure **Authorization Code** and **Refresh Token** are selected.
 - h. Under **Grant type > Advanced > Other grants**, ensure **Implicit (hybrid)** is selected.
2. Add callback URIs for each Region in which your Slack extension will be installed

- a. Compose a callback URI using the following format, replacing *your-region* with your Amazon Quick instance Region for each region where you wish to configure the extension. The Slack extension supports the following Regions: ap-southeast-2, eu-west-1, us-west-2, and us-east-1.

`qbs-cell001.dp.appintegrations.your-region.prod.plato.ai.aws.dev/auth/idc-tti/callback`
 - b. Under **Sign-in redirect URIs**, click on **Add URI** and paste each of the URIs you generated from the previous step.
3. Provide your organization access to the app:
 - a. Under **Assignments** > **Controlled access**, select the groups in your organization that need to have access.
 - b. Under **Assignments** > **Enable immediate access**, select **Enable immediate access with Federation Broker Mode**.
 - c. Click on **Save**.
 4. Note down the **Client ID** and **Client Secret** for the app integration you just created. You will need this in the next steps.

To configure a Trusted Token Issuer

1. Go to your AWS account and navigate to your IAM Identity Center instance.
2. Navigate to **Settings** > **Authentication**.
3. Choose **Create trusted token issuer**.
4. Add the issuer URL, which should follow this template, where *yourOktaDomain* refers to the okta URL for your organization, which may look like `your-organization.okta.com`:

`https://{yourOktaDomain}/oauth2/default`

Note

The issuer URL should be the OIDC discovery endpoint of your identity without the well-known document URI path. If you include the well-known document URI path, this will not work. See Trusted token issuer configuration settings.

5. Choose **Email** as the Identity Provider attribute and IAM Identity Center attribute.
6. Note the trusted token issuer ARN. You will need this in a later step.

After completing these Okta-specific steps, proceed to the [Complete AWS Configuration \(all providers\)](#) section below.

Complete AWS Configuration (all providers)

Follow these steps to set up permissions on AWS Console:

To set up permissions

1. Navigate to Secrets Manager on AWS console.
2. Choose **Store a new secret**.
3. Choose **Other type of secret** and choose the **Plaintext** tab.
4. Your secret should be in the following format and use the app registration client ID and app registration client secret that you saved from the earlier steps:

```
{
  "client_id": "Your app registration client ID",
  "client_secret": "Your app registration client secret"
}
```

5. Navigate to the secret you just created and save the ARN for later.
6. Now navigate to IAM on AWS console.
7. Choose **Access Management, Roles** in the left navigation bar.
8. Choose **Create role**.
9. Choose **Custom trust policy**.
10. Configure the role to trust our service principal for the relevant Region that you selected when configuring your identity provider app integration by adding the following statement replacing *your-region* with the Region you chose when creating your identity provider app integration:

```
{
  "Version": "2012-10-17",
  "Statement": [
```

```
{
  "Effect": "Allow",
  "Principal": {
    "Service": "your-region.prod.appintegrations.plato.aws.internal"
  },
  "Action": "sts:AssumeRole",
  "Condition": {}
}
```

11. Choose **Next**.
12. Provide a name and description and choose **Create role**.
13. Navigate to the role you just created and choose it.
14. Choose **Add Permissions, Create inline policy**.
15. Choose **JSON**.
16. Configure the role with permissions to read secrets from Secrets Manager and permissions to invoke `sso:DescribeTrustedTokenIssuer` by adding the following statement:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "BasePermissions",
      "Effect": "Allow",
      "Action": [
        "secretsmanager:GetSecretValue",
        "sso:DescribeTrustedTokenIssuer"
      ],
      "Resource": "*"
    }
  ]
}
```

17. Provide a policy name and choose **Create policy**.
18. Copy and save the ARN of the IAM role you created for later. You will need it in the next steps.

Now you can follow these steps to create a new extension access configuration that will allow Amazon Quick to integrate with your Slack environment:

To create an extension access configuration

1. Sign in to the Amazon Quick console.
2. Choose the profile picture icon.
3. From the drop-down menu, choose **Manage account**.
4. Under **Permissions**, choose **Extension access**.
5. Choose **New extension access**.
6. If this is your first time setting up extension access, you are prompted to complete the **Trusted Token Issuer Setup**. Configure the following fields and then choose **Next**:
 - **Trusted Token Issuer ARN** – Enter the trusted token issuer ARN that you noted from the earlier steps.
 - **Aud claim** – Enter the client ID from your app registration (Entra ID) or app integration (Okta) that you saved from the earlier steps.

Note

This is a one-time setup that establishes a trusted identity source for all extensions. Once completed, you won't need to do this again for other extension accesses.

7. Select **Slack**. Then, choose **Next**.
8. Configure the following fields:
 - **Name** - A name for your extension is pre-filled for you. You can edit this and enter a descriptive name for the Slack extension (maximum 512 alphanumeric characters, hyphens allowed but no spaces).
 - **Description (optional)** - A description for your extension is pre-filled for you. You can edit this and enter a new description to provide additional context about this extension configuration (maximum 1000 characters).
 - **Slack Workspace ID** - Enter your Slack workspace identifier. Workspace IDs must start with 'T' and be between 1 and 256 alphanumeric characters.
 - **Secrets Role ARN** - Paste the ARN of the IAM role you created from the previous steps.
 - **Secrets ARN** - Paste the ARN of the Secrets Manager secret you created from the previous steps.
9. Choose **Add** to save the new access configuration.

A success message will open up on the top right of your screen.

10. From the success message, choose **View extensions** to finish installing your extension.

 Note

You can also navigate to the installation screen from **Connections > Extensions** in the Amazon Quick menu.

Once created, this extension access configuration enables authors and other admin in your organization to create and deploy Amazon Quick extensions within your Slack environment.

 Note

For your end users to begin using your Slack extension, an admin or author must finish deploying a extension after you configure extension access. Notify your authors that they can view, edit, and complete installation of this extension under **Extensions** in the left navigation once it has been shared. To learn how to do this see Installing your Slack extension in the Slack extension author guide.

Add Slack extension access for accounts using other authentication methods

Follow these steps to create a new extension access configuration that will allow Amazon Quick to integrate with your Slack environment.

1. Sign in to the Amazon Quick console.
2. Choose the profile picture icon.
3. From the drop-down menu, select **Manage account**.
4. Under **Permissions**, select **Extension access**.
5. Choose **New extension access**.
6. Select Slack, then, **Next**.
7. Configure the following fields:

- **Name** - A name for your extension is pre-filled for you. You can edit this and enter a descriptive name for the Slack extension (maximum 512 alphanumeric characters, hyphens allowed but no spaces).
 - **Description** (optional) - A description for your extension is pre-filled for you. You can edit this and enter a new description to provide additional context about this extension configuration (maximum 1000 characters).
 - **Slack Workspace ID** - Enter your Slack workspace identifier. Workspace ID must start with T and be between 1 and 256 alphanumeric characters long.
8. Select **Add** to save the new access configuration.

A success message will open up on the top right of your screen.

9. From the success message, select **View extensions** to finish installing your extension.

 Note

You can also navigate to the installation screen from **Connections > Extensions** in the Amazon Quick menu.

Once created, this extension access configuration enables authors and other admins in your organization to deploy Amazon Quick Slack extensions in their workspace.

 Note

For your end users to begin using your Slack extension, an admin or author must finish deploying a extension after you configure extension access. Notify your authors that they can view, edit, and complete installation of this extension under **Extensions** in the left navigation once it has been shared. To learn how to do this see [Installing your Slack extension in the Slack extension author guide](#).

Edit Slack extension access

Use these steps to modify the configuration settings of an existing Slack extension access.

1. Sign in to the Amazon Quick console.
2. Choose the profile picture icon.

3. From the drop-down menu, select **Manage account**.
4. Under **Permissions**, select **Extension access**.
5. Select the three dot menu icon for the Slack Extension you need to edit.
6. Select **Edit**.
7. Edit the configuration as required and select **Save** to confirm the changes.

Your changes to the Slack extension access configuration are saved and will take effect immediately.

Delete Slack extension access

Follow these steps to permanently remove a Slack extension access configuration. This action cannot be undone.

1. Sign in to the Amazon Quick console.
2. Choose the profile picture icon.
3. From the drop-down menu, select **Manage account**.
4. Under **Permissions**, select **Extension access**.
5. Select the three dot menu icon for the Slack Extension you need to delete.
6. Select **Delete**.
7. Enter the word, "confirm", and select **DELETE**.

Note

Deleting a extension access removes access for all users in your Slack workspace and deletes all extensions created for Slack. If delete extension access fails, the admin must switch to the author view and delete the Slack extensions that are using the configured extension access before returning to delete the extension access.

With Slack extension access configured, your team can now use **@Amazon Quick** mentions in channels and direct messages to access AI assistance and organizational knowledge directly within their Slack workspace.

Amazon Quick Microsoft Teams extension

The Quick extension for Microsoft Teams integrates AI-powered assistance directly into your team communication workflows. Users can mention @Amazon Quick in conversations to access company knowledge, use configured action connectors, and get contextual assistance without leaving their Teams environment.

The Teams extension enables users to:

- Mention **@Amazon Quick** in conversations in Teams channels to add it as a collaborator.
- Use actions from connectors configured in Amazon Quick.
- Access any company knowledge sources added to your Amazon Quick instance from within Microsoft Teams.

Important

- The Amazon Quick customer integrating Microsoft Teams must have a paid Microsoft Teams organization with an *M365 subscription for their organization*.
- Amazon Quick does not use your user data for service improvement or for training its underlying large language models (LLMs).

Topics

- [Prerequisites for Microsoft Teams extension](#)
- [Microsoft Teams extension permissions](#)
- [Configure Microsoft Teams extension access](#)

Prerequisites for Microsoft Teams extension

Before configuring access for the Amazon Quick extension to Microsoft Teams, administrators must complete the following steps:

1. Have a Microsoft 365 Business subscription and be a Global Admin or have administrative permissions (specifically `AppCatalog.ReadWrite.All`).
2. Have a Amazon Quick instance.

3. Your Microsoft 365 tenant ID. You can find this by going to the Azure portal > **Azure Active Directory** > **Properties**, or by using PowerShell. For detailed steps, see [How to find your tenant ID - Microsoft Entra](#) in the Microsoft Learn portal.

If you configured the authentication to connect to Amazon Quick with IAM Identity Center, complete the following additional steps:

1. Ensure you have an IAM Identity Center instance enabled on your AWS account.
2. If you are connecting an external IAM provider to IAM Identity Center, ensure that every user under your IAM provider configuration has an email associated with them.
3. If you are using Entra ID, set up SCIM identity propagation between the Microsoft Entra ID instance and IAM Identity Center. For detailed steps, see [Configure SAML and SCIM with Microsoft Entra ID and IAM Identity Center](#).

Microsoft Teams extension permissions

The Amazon Quick Microsoft Teams extension uses Microsoft's Resource-Specific Consent (RSC) framework. The bot only gets access to the specific meetings and chats where it is installed — not tenant-wide access to all Teams data. These permissions are granted automatically when the app is added to a meeting or chat; no additional admin action is needed.

The following table lists the RSC permissions used by the Teams extension.

Permission	Purpose	Feature
ChatMessage.Read.Chat	Read chat messages for AI assistant context	General chat assistant
ChatMember.Read.Chat	Resolve attendee emails and identify the meeting organizer for summary delivery	Meeting Summary
OnlineMeeting.ReadBasic.Chat	Read meeting name, schedule, and organizer info	Meeting Summary
OnlineMeetingTranscript.Read.Chat	Retrieve transcripts after a meeting ends	Meeting Summary

Permission	Purpose	Feature
OnlineMeetingParticipant.Read.Chat	Read participant join/leave times and roles	Meeting Summary
ChannelMessage.Read.Group	Read thread context when the bot is @mentioned in a channel	General chat assistant
ChannelMeeting.ReadBasic.Group	Read basic properties of channel meetings	Meeting Summary
ChannelMeetingTranscript.Read.Group	Retrieve transcripts of channel meetings	Meeting Summary
ChannelMeetingParticipant.Read.Group	Read participant info for channel meetings	Meeting Summary

Installation permission (revoke after install)

The following permission is automatically granted during app installation and must be revoked afterward. It is not required for ongoing operation.

Permission	Purpose	Required admin action
AppCatalog.ReadWrite.All	Read and write to all app catalogs (used during installation only)	Revoke after installation via Microsoft Entra > Enterprise Applications > Permissions

Important

After installation, go to Microsoft Entra > **Enterprise Applications**, find the app starting with "qbs" and ending with "teams", navigate to **Permissions**, and revoke AppCatalog.ReadWrite.All.

Where to review permissions

Use the following locations to review permissions for the Microsoft Teams extension:

- **RSC Permissions:** Teams Admin Center > **Teams apps** > **Manage apps** > **Amazon Quick** > **Permissions**
- **Application Permissions:** Microsoft Entra > **Enterprise Applications** > find the app starting with "qbs" and ending with "teams" > **Permissions**

Configure Microsoft Teams extension access

As an administrator, you must allow the Amazon Quick Microsoft Teams to connect to your Amazon Quick application environment. You can use the Amazon Quick console to manage extension access configurations.

User attribute mapping

When you configure a Microsoft Teams extension, user identity is mapped by default using the following attributes:

- **Amazon Quick user attribute** - Email address is used to map Amazon Quick users to their corresponding Microsoft 365 accounts. The system uses the email address to establish the connection between user identities.
- **M365 Office add-in user attribute** - User Principal Name (UPN) is used to match against Microsoft 365 user accounts. Users need to use the User Principal Name (UPN) to sign in to Microsoft 365. Email address that works on both cloud and on-premise systems.

These default mappings ensure secure and accurate user identification across both platforms without requiring additional configuration.

Topics

- [Add Microsoft Teams extension access for accounts using IAM Identity Center](#)
- [Add Microsoft Teams extension access for accounts using other authentication methods](#)
- [Edit Microsoft Teams extension access](#)
- [Delete Microsoft Teams extension access](#)

Add Microsoft Teams extension access for accounts using IAM Identity Center

Configuring extension access with IAM Identity Center requires completing steps specific to your identity provider (for example, Entra ID or Okta) followed by common setup steps in AWS.

Configure IAM Identity Center with Entra ID

Follow these steps only if you are using IAM Identity Center with Entra ID to set up and configure an Azure tenant on your Microsoft Azure portal:

To set up an Azure tenant

1. In the Azure account, create a new app registration.
 - a. Go to **App registrations**.
 - b. In the **App registrations** screen, choose **New registration**. Under the **Supported account types** option, choose **Accounts in this organizational directory only (Personal use only - Single tenant)**. Once finished, choose **Register**.
 - c. Note the client ID. You will need this later.
 - d. Create a client secret for the app registration. Copy and save the secret's **Value** attribute. You need it later. This value is only shown once and cannot be retrieved after you leave the page.
2. Add callback URLs for each Region in which your Teams extension will be installed.
 - a. Navigate to the app registration's **Authentication** tab.
 - b. Choose **Platform Configurations, Add a platform**.
 - c. Choose **Web**.
 - d. Compose a callback URL using the following format, replacing *your-region* with your Amazon Quick instance Region. The Teams extension supports the following Regions: ap-southeast-2, eu-west-1, us-west-2, and us-east-1.

`qbs-cell001.dp.appintegrations.your-region.prod.plato.ai.aws.dev/auth/idc-tti/callback`
 - e. Insert the callback URL as the redirect URI and choose **Configure** when done.

Follow these steps to configure a Trusted Token Issuer on your IAM Identity Center instance in your AWS Console:

To configure a Trusted Token Issuer

1. Go to your AWS account and navigate to your IAM Identity Center instance.
2. Navigate to **Settings, Authentication**.
3. Choose **Create trusted token issuer**.
4. Add the issuer URL, which should follow this template, where *Tenant ID* refers to your Entra tenant ID:

```
login.microsoftonline.com/Tenant ID/v2.0
```

Note

The issuer URL should be the OIDC discovery endpoint of your identity without the well-known document URI path. If you include the well-known document URI path, this will not work. See Trusted token issuer configuration settings.

5. Choose **Email** as the Identity Provider attribute and IAM Identity Center attribute.
6. Note the trusted token issuer ARN. You will need this in a later step.

After completing these Entra ID-specific steps, proceed to the [Complete AWS Configuration \(all providers\)](#) section below.

Configure IAM Identity Center with Okta

Follow these steps only if you are using IAM Identity Center with Okta to set up and configure your App Integration in the Okta Admin console:

To set up an Okta Application

1. In your Okta account, create a new Okta App Integration.
 - a. In your Okta Admin console, navigate to **Applications > Applications**.
 - b. Click on **Create App Integration**.
 - c. For the Sign-in method, select **OIDC - OpenID Connect**.
 - d. For the Application type, select **Web Application**.
 - e. Click on **Next**.
 - f. Provide an App integration name.

- g. Under **Grant type** > **Core grants**, ensure **Authorization Code** and **Refresh Token** are selected.
 - h. Under **Grant type** > **Advanced** > **Other grants**, ensure **Implicit (hybrid)** is selected.
 2. Add callback URIs for each Region in which your Teams extension will be installed
 - a. Compose a callback URI using the following format, replacing *your-region* with your Amazon Quick instance Region for each region where you wish to configure the extension. The Teams extension supports the following Regions: ap-southeast-2, eu-west-1, us-west-2, and us-east-1.

```
qbs-cell001.dp.appintegrations.your-region.prod.plato.ai.aws.dev/auth/idc-tti/callback
```
 - b. Under **Sign-in redirect URIs**, click on **Add URI** and paste each of the URIs you generated from the previous step.
3. Provide your organization access to the app:
 - a. Under **Assignments** > **Controlled access**, select the groups in your organization that need to have access.
 - b. Under **Assignments** > **Enable immediate access**, select **Enable immediate access with Federation Broker Mode**.
 - c. Click on **Save**.
4. Note down the **Client ID** and **Client Secret** for the app integration you just created. You will need this in the next steps.

To configure a Trusted Token Issuer

1. Go to your AWS account and navigate to your IAM Identity Center instance.
2. Navigate to **Settings** > **Authentication**.
3. Choose **Create trusted token issuer**.
4. Add the issuer URL, which should follow this template, where *yourOktaDomain* refers to the okta URL for your organization, which may look like `your-organization.okta.com`:

```
https://{yourOktaDomain}/oauth2/default
```

Note

The issuer URL should be the OIDC discovery endpoint of your identity without the well-known document URI path. If you include the well-known document URI path, this will not work. See Trusted token issuer configuration settings.

5. Choose **Email** as the Identity Provider attribute and IAM Identity Center attribute.
6. Note the trusted token issuer ARN. You will need this in a later step.

After completing these Okta-specific steps, proceed to the [Complete AWS Configuration \(all providers\)](#) section below.

Complete AWS Configuration (all providers)

Follow these steps to set up permissions on AWS Console:

To set up permissions

1. Navigate to Secrets Manager on AWS console.
2. Choose **Store a new secret**.
3. Choose **Other type of secret** and choose the **Plaintext** tab.
4. Your secret should be in the following format and use the app registration client ID and app registration client secret that you saved from the earlier steps:

```
{
  "client_id": "Your app registration client ID",
  "client_secret": "Your app registration client secret value"
}
```

5. Navigate to the secret you just created and save the ARN for later.
6. Now navigate to IAM on AWS console.
7. Choose **Access Management, Roles** in the left navigation bar.
8. Choose **Create role**.
9. Choose **Custom trust policy**.
10. Configure the role to trust our service principal for the relevant Region that you selected when configuring your identity provider app integration by adding the following statement

replacing *your-region* with the Region you chose when creating your identity provider app integration:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": "your-region.prod.appintegrations.plato.aws.internal"
      },
      "Action": "sts:AssumeRole",
      "Condition": {}
    }
  ]
}
```

11. Choose **Next**.
12. Provide a name and description and choose **Create role**.
13. Navigate to the role you just created and choose it.
14. Choose **Add Permissions, Create inline policy**.
15. Choose **JSON**.
16. Configure the role with permissions to read secrets from Secrets Manager and permissions to invoke `sso:DescribeTrustedTokenIssuer` by adding the following statement:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "BasePermissions",
      "Effect": "Allow",
      "Action": [
        "secretsmanager:GetSecretValue",
        "sso:DescribeTrustedTokenIssuer"
      ],
      "Resource": "*"
    }
  ]
}
```

```
}
```

17. Provide a policy name and choose **Create policy**.
18. Copy and save the ARN of the IAM role you created for later. You will need it in the next steps.

Now you can follow these steps to create a new extension access configuration that will allow Amazon Quick to integrate with your Microsoft Teams environment:

To create an extension access configuration

1. Sign in to the Amazon Quick console.
2. Choose the profile picture icon.
3. From the drop-down menu, choose **Manage account**.
4. Under **Permissions**, choose **Extension access**.
5. Choose **New extension access**.
6. If this is your first time setting up extension access, you are prompted to complete the **Trusted Token Issuer Setup**. Configure the following fields and then choose **Next**:
 - **Trusted Token Issuer ARN** – Enter the trusted token issuer ARN that you noted from the earlier steps.
 - **Aud claim** – Enter the client ID from your app registration (Entra ID) or app integration (Okta) that you saved from the earlier steps.

Note

This is a one-time setup that establishes a trusted identity source for all extensions. Once completed, you won't need to do this again for other extension accesses.

7. Select **Microsoft Teams**. Then, choose **Next**.
8. Configure the following fields:
 - **Name** - A name for your extension is pre-filled for you. You can edit this and enter a descriptive name for the Teams extension (maximum 512 alphanumeric characters, hyphens allowed but no spaces).

- **Description (optional)** - A description for your extension is pre-filled for you. You can edit this and enter a new description to provide additional context about this extension configuration (maximum 1000 characters).
 - **M365 tenant ID** - Enter your Microsoft 365 tenant identifier (must be 36 characters).
 - **Secrets Role ARN** - Paste the ARN of the IAM role you created from the previous steps.
 - **Secrets ARN** - Paste the ARN of the Secrets Manager secret you created from the previous steps.
9. Choose **Add** to save the new access configuration.

A success message will open up on the top right of your screen.

10. From the success message, choose **View extensions** to finish installing your extension.

 Note

You can also navigate to the installation screen from **Connections > Extensions** in the Amazon Quick menu.

Once created, this extension access configuration enables authors and other admin in your organization to create and deploy Amazon Quick extensions within your Microsoft Teams environment.

 Note

For your end users to begin using your Microsoft Teams extension, an admin or author must finish deploying a extension after you configure extension access. Notify your authors that they can view, edit, and complete installation of this extension under **Extensions** in the left navigation once it has been shared. To learn how to do this see Installing your Microsoft Teams extension in the Microsoft Teams extension author guide.

Add Microsoft Teams extension access for accounts using other authentication methods

Follow these steps to create a new extension access configuration that will allow Amazon Quick to integrate with your Microsoft Teams environment.

1. Sign in to the Amazon Quick console.

2. Choose the profile picture icon.
3. From the drop-down menu, select **Manage account**.
4. Under **Permissions**, select **Extension access**.
5. Choose **New extension access**.
6. Select **Microsoft Teams**, then, **Next**.
7. Configure the following fields:
 - **Name** - A name for your extension is pre-filled for you. You can edit this and enter a descriptive name for the Teams extension (maximum 512 alphanumeric characters, hyphens allowed but no spaces).
 - **Description** (optional) - A description for your extension is pre-filled for you. You can edit this and enter a new description to provide additional context about this extension configuration (maximum 1000 characters).
 - **M365 tenant ID** - Enter your Microsoft 365 tenant identifier (must be 36 characters).
8. Select **Add** to save the new access configuration.

A success message will open up on the top right of your screen.

9. From the success message, select **View extensions** to finish installing your extension.

 Note

You can also navigate to the installation screen from **Connections > Extensions** in the Amazon Quick menu.

Once created, this extension access configuration enables authors and other admins in your organization to deploy your Microsoft Teams Amazon Quick extension within your Microsoft Teams environment.

 Note

For your end users to begin using your Microsoft Teams extension, an admin or author must finish deploying a extension after you configure extension access. Notify your authors that they can view, edit, and complete installation of this extension under **Extensions** in the left

navigation once it has been shared. To learn how to do this see [Installing your Microsoft Teams extension in the Microsoft Teams extension author guide](#).

Edit Microsoft Teams extension access

Use these steps to modify the configuration settings of an existing Microsoft Teams extension access.

1. Sign in to the Amazon Quick console.
2. Choose the profile picture icon.
3. From the drop-down menu, select **Manage account**.
4. Under **Permissions**, select **Extension access**.
5. Select the three dot menu icon for the Microsoft Teams extension you need to edit.
6. Select **Edit**.
7. Edit the configuration as required and select **Save** to confirm the changes.

Your changes to the Microsoft Teams extension access configuration are saved and will take effect immediately.

Delete Microsoft Teams extension access

Follow these steps to permanently remove a Microsoft Teams extension access configuration. This action cannot be undone.

1. Sign in to the Amazon Quick console.
2. Choose the profile picture icon.
3. From the drop-down menu, select **Manage account**.
4. Under **Permissions**, select **Extension access**.
5. Select the three dot menu icon for the Microsoft Teams Extension you need to delete.
6. Select **Delete**.
7. Enter the word, "confirm", and select **DELETE**.

Note

Deleting a extension access removes access for all users in your M365 tenant and deletes all extensions created for Teams. If delete extension access fails, the admin must switch to the author view and delete the Teams extensions that are using the configured extension access before returning to delete the extension access.

With Microsoft Teams extension access configured, your team can now use **@Quick** mentions in conversations to access AI assistance, company knowledge, and connectors directly within their Teams environment.

Custom permissions

Applies to

Applies to: Enterprise Edition

Intended audience: Administrators and Amazon Quick developers

In Enterprise edition, you can use custom permissions profiles to restrict the capabilities and features that people can access in Amazon Quick. Custom permissions is a scope-down policy. It can only restrict what a user has access to based on the default set of capabilities available to their role. It cannot be used to grant capabilities that a user does not already have.

You can configure custom permissions at the account, role (admin, author, reader), and user levels for all identity types in Quick. You can assign custom permissions profiles to users and roles directly in the Amazon Quick admin console, or by using the Amazon Quick API and AWS Command Line Interface. The admin console also provides a Check permissions feature that lets you verify which profile is active for any user and at which level (user, role, or account) it applies.

Permissions precedence

If a user has custom permissions assigned at multiple levels, Quick evaluates them in the following order of precedence. The most specific level wins.

1. **User-level** – Custom permissions assigned directly to a user take highest priority.
2. **Role-level** – Custom permissions assigned to a role (Admin, Author, Reader) apply to all users in that role unless overridden by a user-level assignment.

3. **Account-level** – Custom permissions assigned at the account level apply to all users unless overridden by a role-level or user-level assignment.

The first match wins. A user with a user-level allow-by-default profile is not subject to an account-level Deny by Default profile. Assign Deny by Default profiles at the appropriate level for your requirements.

 Tip

For an in-depth explanation of using custom permissions to establish enterprise controls in Quick, see [Establishing enterprise governance in Amazon Quick using custom permissions](#) on the AWS Business Intelligence Blog.

The following limitations apply to custom permissions.

- You can't grant permissions that are above a user's default role. For example, if a user has reader access, you can't grant permissions for that user to edit dashboards.
- To customize user or role permissions, you need to be a Quick administrator with the following IAM permissions:
 - `quicksight:CreateCustomPermissions`
 - `quicksight>DeleteCustomPermissions`
 - `quicksight:DescribeCustomPermissions`
 - `quicksight:ListCustomPermissions`
 - `quicksight:UpdateCustomPermissions`
 - `quicksight>DeleteAccountCustomPermission`
 - `quicksight>DeleteRoleCustomPermission`
 - `quicksight>DeleteUserCustomPermission`
 - `quicksight:DescribeAccountCustomPermission`
 - `quicksight:DescribeRoleCustomPermission`
 - `quicksight:ListCustomPermissionAssignments`
 - `quicksight:UpdateAccountCustomPermission`
 - `quicksight:UpdateRoleCustomPermission`
 - `quicksight:UpdateUserCustomPermission`

You can create custom permission profiles to restrict access to any combination of the following features. Parent capabilities let you restrict an entire asset's feature set. When you deny a parent capability, all of its child features are also denied. The reverse isn't true: denying every child feature doesn't deny the parent capability. This is because a parent can control functionality that isn't exposed as a separate child feature. To restrict a capability completely, deny the parent.

 **Parent capability overrides child features**

You can't allow an individual child feature while its parent capability is denied. If you set a parent capability to DENY and one of its child features to ALLOW, the child feature remains restricted. Quick accepts this combination without an error, so verify the result in the console after you apply it. To allow specific child features, leave the parent capability allowed and deny only the child features you want to restrict.

Features that show **Not applicable** in the **Parent capability** column have no parent capability. Restrict them individually, or restrict their whole category. For more information, see [the section called “Deny by Default”](#).

Quick capabilities and features

The following tables list the capabilities and features that you can restrict with custom permissions profiles. Capabilities are organized into four groups that match the console layout.

The **Badge** column indicates the governance category for each capability. **AI** means the capability belongs to the AI governance category, which is the category that you can restrict wholesale with Deny by Default. **Contains AI** means the parent capability is not itself in the AI category but has children that are.

AI capabilities

Use this group to control access to AI-powered features such as natural language Q&A, executive summaries, anomaly detection, and AI-generated insights.

Feature	Amazon Quick behavior	Badge	Parent capability
Apps	Restricts all app capabilities, including creating, updating, sharing, and using apps.	AI	Not applicable

Feature	Amazon Quick behavior	Badge	Parent capability
Access apps native data store when using existing apps	Restricts apps from accessing their native data store when users run existing apps.	AI	Apps
Create and update apps	Users can't create or update apps.	AI	Apps
Invoke AI inference when using existing apps	Restricts apps from invoking AI inference when users run existing apps.	AI	Apps
Share apps	Users can't share apps with other users. Apps that were previously shared remain accessible to those users.	AI	Apps
Automate	Restricts access to Automate. Users can't create, update, or run automations.	AI	Not applicable
Chat agents	Users can't view or access any chat agents. The agent library and navigation are hidden.	AI	Not applicable
Create chat agents	Users can't create new chat agents.	AI	Chat agents
Share chat agents	Users can't share chat agents with other users. Chat agents that were previously shared remain accessible to those users.	AI	Chat agents
Extensions	Restricts access to Quick through all browser and Microsoft Office extensions.	AI	Not applicable

Feature	Amazon Quick behavior	Badge	Parent capability
Browser extension	Prevents access to Quick through the browser extension for Chrome, Firefox, and Edge.	AI	Extensions
Excel extension	Prevents access to Quick through the Microsoft Excel add-in.	AI	Extensions
Outlook extension	Prevents access to Quick through the Microsoft Outlook add-in.	AI	Extensions
PowerPoint extension	Prevents access to Quick through the Microsoft PowerPoint add-in.	AI	Extensions
Word extension	Prevents access to Quick through the Microsoft Word add-in.	AI	Extensions
Flows	Restricts all Flows capabilities, including creating, sharing, and running flows.	AI	Not applicable
(Preview) Enable UI agent to perform browser tasks	Restricts the Flows UI agent from performing browser tasks.	AI	Flows
All eligible users can review and approve flow sharing requests	Restricts which users can review and approve flow sharing requests.	AI	Flows
Allow creators to share without approval	Creators can't share flows without approval.	AI	Flows
Enable Bedrock model usage in General knowledge step for output refinement	Restricts usage of Bedrock models.	AI	Flows

Feature	Amazon Quick behavior	Badge	Parent capability
Quick desktop	Restricts all Quick desktop capabilities.	AI	Not applicable
Automate browser actions	Restricts Quick desktop from automating browser actions.	AI	Quick desktop
Build and manage engrams	Users can't build or manage engrams in Quick desktop.	AI	Quick desktop
Create and access knowledge memory	Users can't create or access knowledge memory in Quick desktop.	AI	Quick desktop
Execute code locally	Restricts Quick desktop from running code locally.	AI	Quick desktop
Generate images with AI	Users can't generate images with AI in Quick desktop.	AI	Quick desktop
Manage agents	Users can't manage agents in Quick desktop.	AI	Quick desktop
Manage tasks	Users can't manage tasks in Quick desktop.	AI	Quick desktop
Receive chat notifications	Restricts chat notifications in Quick desktop.	AI	Quick desktop
Use coding agents (ACP)	Users can't use coding agents (ACP) in Quick desktop.	AI	Quick desktop
Research	Restricts access to Research.	AI	Not applicable
Skills	Restricts access to skills, including sharing skills.	AI	Not applicable
Share skills with individuals	Users can't share skills with individuals.	AI	Skills

Feature	Amazon Quick behavior	Badge	Parent capability
Spaces	Restricts all Spaces capabilities, including creating and sharing spaces.	AI	Not applicable
Create spaces	Users can't create new spaces. Existing spaces remain accessible and unchanged.	AI	Spaces
Share spaces	Users can't share spaces with other users. Spaces that were previously shared remain accessible to those users.	AI	Spaces
Triggers	Restricts all trigger capabilities.	AI	Not applicable
Inbound email triggers	Users can't create or use inbound email triggers.	AI	Triggers
Quick event triggers	Users can't create or use Quick event triggers.	AI	Triggers
Schedule triggers	Users can't create or use schedule triggers.	AI	Triggers
Use internet to enhance results	Restricts usage of web-based search in Chat agents and Research.	AI	Not applicable

BI capabilities

Use this group to control creation and management of core business intelligence assets: dashboards, analyses, datasets, SPICE capacity, and scheduled reports.

Feature	Amazon Quick behavior	Badge	Parent capability
Analyses	Restricts all analysis capabilities.	Contains AI	Not applicable
Build calculation with AI	Users can't build calculated fields with AI.	AI	Analyses
Edit visual with AI	Users can't edit visuals with AI.	AI	Analyses
Generate analyses	Users can't generate analyses with AI.	AI	Analyses
Sharing analyses	Access to the Share option on the File menu is disabled for analyses.	Not applicable	Analyses
Adding or running anomaly detection	Access to the Add anomaly to sheet option on the Insights menu is disabled for analyses. Users can't add anomaly detection to sheets.	AI	Not applicable
Content within scheduled email reports	Restricts including content within scheduled email reports.	Not applicable	Not applicable
Creating or updating scheduled email reports	Restricts creating or updating scheduled email reports.	Not applicable	Not applicable
Creating or updating themes	Users can't create new custom themes or edit existing themes. Users can still view or apply existing themes.	Not applicable	Not applicable
CSV attachments in scheduled email reports	Restricts including CSV attachments in scheduled email reports.	Not applicable	Not applicable

Feature	Amazon Quick behavior	Badge	Parent capability
Dashboards	Restricts all dashboard capabilities.	Contains AI	Not applicable
Create executive summary	Users can't create executive summaries with AI for dashboards.	AI	Dashboards
Sharing dashboards	Access to the share icon on the navigation menu is disabled for dashboards.	Not applicable	Dashboards
Subscribing to scheduled email reports	Users can't subscribe to scheduled email reports.	Not applicable	Dashboards
Creating or updating all datasets	Access to creating or updating all datasets is disabled.	Not applicable	Not applicable
Creating or updating only SPICE datasets	Access to creating or updating SPICE datasets is disabled.	Not applicable	Creating or updating all datasets
Sharing datasets	Access to sharing datasets is disabled.	Not applicable	Not applicable
Excel attachments in scheduled email reports	Restricts including Excel attachments in scheduled email reports.	Not applicable	Not applicable
Export sheet to PDF	Access to the Export to PDF option on the File menu is disabled for analyses. Access to the Generate PDF option on the Export menu is disabled for dashboards.	Not applicable	Not applicable

Feature	Amazon Quick behavior	Badge	Parent capability
Export visual to CSV	Access to the Export to CSV option on the actions menu for each visual is disabled for both analyses and dashboards.	Not applicable	Not applicable
Export visual to Excel	Access to the Export to Excel option on the actions menu for each table is disabled for both analyses and dashboards.	Not applicable	Not applicable
PDF attachments in scheduled email reports	Restricts including PDF attachments in scheduled email reports.	Not applicable	Not applicable
Print sheet	Access to the Print option on the File menu is disabled for analyses. Access to the Print option on the Export menu is disabled for dashboards.	Not applicable	Not applicable
Scenarios	Restricts access to scenarios.	AI	Not applicable
Managing shared folders	Restricts creating, updating, deleting, and viewing shared folders, adding assets to shared folders, and sharing folders. Doesn't prevent inheriting access to assets shared through folder membership.	Not applicable	Not applicable
Creating shared folders	Restricts creating shared folders.	Not applicable	Managing shared folders

Feature	Amazon Quick behavior	Badge	Parent capability
Renaming shared folders	Restricts renaming shared folders.	Not applicable	Managing shared folders
Stories	Restricts access to stories.	AI	Not applicable
Creating or updating threshold alerts	Users can't create or update threshold alerts.	Not applicable	Not applicable
Topics	Restricts access to topics.	AI	Not applicable
Viewing account SPICE capacity	Restricts retrieving the account's SPICE capacity.	Not applicable	Not applicable

Connectors

Use this group to control data source connections, including adding, editing, and removing third-party integrations such as Salesforce, Google Sheets, and databases.

Feature	Amazon Quick behavior	Badge	Parent capability
Actions	Restricts all connector capabilities.	AI	Not applicable
Creating or updating all data sources	Access to creating or updating all data sources is disabled.	Not applicable	Not applicable
Sharing data sources	Access to sharing data sources is disabled.	Not applicable	Not applicable
Knowledge base	Restricts all knowledge base capabilities.	AI	Not applicable

Feature	Amazon Quick behavior	Badge	Parent capability
Creating or updating all knowledge bases	Users can't create or update knowledge bases.	AI	Knowledge base
Share all knowledge bases	Users can't share knowledge bases with other users.	AI	Knowledge base

Connector features

In addition to the features listed in the preceding sections, you can restrict access to individual connectors. Each connector supports the following permissions:

- **Create and Update action** – Restricts the ability to create or update actions for the connector.
- **Share action** – Restricts the ability to share actions for the connector.
- **Use action** – Restricts the ability to use actions for the connector.

These permissions are available under **Actions**, in the **Connectors** group of the **Capabilities & features** section. For a list of available connectors, see [Connectors](#).

Knowledge base connectors

You can restrict access to individual knowledge base connectors. Each knowledge base connector supports the following permissions:

- **Create and Update** – Restricts the ability to create or update knowledge bases for the connector.
- **Share** – Restricts the ability to share knowledge bases for the connector.
- **Use** – Restricts the ability to use knowledge bases for the connector.

These permissions are available under **Knowledge base**, in the **Connectors** group of the **Capabilities & features** section. For a list of available knowledge base connectors, see [Knowledge bases](#).

Administration

Use this group to control account-level settings, including user management, permission sets, billing, audit logs, and governance policies.

Feature	Amazon Quick behavior	Badge	Parent capability
Allow users to upgrade or request upgrades	Users can't upgrade their own role or request a role upgrade.	Not applicable	Not applicable

Creating custom permissions profiles

Custom permissions profiles can be created for Amazon Quick accounts that are integrated with IAM Identity Center, Active Directory, or for Quick accounts that have Quick managed users. The identity type that an Quick account uses determines the way an Quick admin configures a custom permissions profile.

Creating or editing a custom permissions profile (Quick console)

To create a custom permissions profile

1. Open the [Quick console](#).
2. Choose **Manage Quick**.
3. In the left navigation, choose **Permissions**, and then choose **Custom permissions**.
4. On the **Custom permissions** page, choose **Create profile**.
5. For **Profile name**, enter a descriptive name for the profile.
6. (Optional) For **Description**, enter a description.
7. (Optional) To restrict a whole capability category, turn on the corresponding toggle in the **Restrict capabilities** section – for example, **Restrict AI capabilities**. Quick restricts all capabilities in that category for users assigned to this profile, including capabilities that it launches later. Allow any capabilities in the category that your users need in the **Capabilities & features** section. For more information, see [the section called “Deny by Default”](#).
8. In the **Capabilities & features** section, choose which capabilities to restrict. Capabilities are organized into four groups: **AI capabilities**, **BI capabilities**, **Connectors**, and **Administration**. Expand a capability to view and configure its individual features. Review the preview panel on the right to see the effect of your selections.
 - To restrict an entire group of features, select the parent capability.
 - To restrict specific features, expand the parent capability and select only the features you want to restrict.

9. Choose **Create**.

Important

Test custom permissions configurations in a development or staging account before applying them to production users to verify the behavior matches your expectations.

After you create the profile, Quick takes you to the assignment page where you can assign it to users, roles, or the account. For more information, see [the section called “Assign profiles”](#).

To edit an existing custom permissions profile

1. On the **Custom permissions** page, locate the profile you want to edit.
2. Choose the actions menu next to the profile, and then choose **Edit**.
3. Make your changes to the profile configuration, including the profile name, Deny by Default settings, and capability restrictions.
4. Choose **Update** to save your changes.

Creating a custom permissions profile (AWS CLI)

Before you begin, you need to set up and configure the AWS CLI. For more information about installing the AWS CLI, see [Install or update the latest version of the AWS CLI](#) and [Configure the AWS CLI](#) in the AWS Command Line Interface User guide.

Use the latest CLI version

Use the latest version of the AWS Command Line Interface. Quick adds capability identifiers to the AWS Command Line Interface as it releases new capabilities. If you receive Unknown parameter in Capabilities for a capability that appears in the console, update the AWS Command Line Interface and try again.

The following example creates a custom permissions profile with ExportToCsv and ExportToPdf denied.

```
aws quicksight create-custom-permissions \
```

```
--aws-account-id AWSACCOUNTID \  
--custom-permissions-name PERMISSIONNAME \  
--capabilities '{"ExportToCsv": "DENY", "ExportToPdf": "DENY"}'
```

Assigning custom permissions profiles

After you create a custom permissions profile, you can assign it to users, roles, or accounts.

Users with sufficient permissions can also use the [AWS::QuickSight::CustomPermissions](#) CloudFormation resource to manage Amazon Quick custom permissions profiles.

Assigning a custom permissions profile (Quick console)

You can assign a custom permissions profile to specific users, roles, or to the entire account.

Assigning a profile replaces any existing assignment at that level. A user, role, or account can only have one active custom permissions profile at a time.

Note

If a user or role already has a custom permissions profile assigned, the new assignment replaces the previous one. The original profile is no longer active for that user or role. Similarly, if an account-level profile is already active, Quick displays a warning before you replace it.

To assign a custom permissions profile

1. Open the [Quick console](#).
2. Choose **Manage Quick**.
3. In the left navigation, choose **Permissions**, and then choose **Custom permissions**.
4. On the **Custom permissions** page, locate the profile you want to assign. Choose the actions menu next to the profile, and then choose **Assign**.

Tip

You can also assign a profile immediately after creating one. After you choose **Create** on the profile creation page, Quick takes you directly to the assignment page.

5. On the assignment page, configure who receives this profile:

- **Users** – Enter the usernames of the users you want to assign this profile to. You can add multiple users.
 - **Roles** – Select the Quick roles (Admin, Author, Reader) that you want to assign this profile to.
 - **Set as account default** – Select this option to apply the profile as the account-level default. If an account-level profile is already active, Quick displays a warning indicating that the existing assignment will be replaced.
6. When you are ready, choose **Assign**. You are returned to the custom permissions list page.

Assigning custom permissions profiles (AWS CLI)

Before you begin, you need to set up and configure the AWS CLI. For more information about installing the AWS CLI, see [Install or update the latest version of the AWS CLI](#) and [Configure the AWS CLI](#) in the AWS Command Line Interface User guide. You also need permissions to use the Quick API.

Updating custom permissions assignments

The following example updates the custom permissions that are assigned to a role.

```
aws quicksight update-role-custom-permission \  
--role ROLE \  
--aws-account-id AWSACCOUNTID \  
--namespace default \  
--custom-permissions-name PERMISSIONNAME
```

The following example applies a custom permissions profile to a user.

```
aws quicksight update-user-custom-permission \  
--aws-account-id AWSACCOUNTID \  
--namespace default \  
--user-name USER_NAME \  
--custom-permissions-name PERMISSIONNAME
```

The following example updates the custom permissions that are assigned to an account.

```
aws quicksight update-account-custom-permission \  
--aws-account-id AWSACCOUNTID \  
--custom-permissions-name PERMISSIONNAME
```

Removing custom permissions assignments

The following example deletes a custom permissions profile from a role.

```
aws quicksight delete-role-custom-permission \  
--role ROLE \  
--aws-account-id AWSACCOUNTID \  
--namespace default
```

The following example deletes a custom permissions profile from a user.

```
aws quicksight delete-user-custom-permission \  
--user-name USER_NAME \  
--aws-account-id AWSACCOUNTID \  
--namespace default
```

The following example removes a custom permissions profile from an account.

```
aws quicksight delete-account-custom-permission \  
--aws-account-id AWSACCOUNTID
```

Describing custom permissions assignments

The following example returns the custom permissions profile that is assigned to a role.

```
aws quicksight describe-role-custom-permission \  
--role ROLE \  
--aws-account-id AWSACCOUNTID \  
--namespace default
```

The following example returns the custom permissions profile that is assigned to an account.

```
aws quicksight describe-account-custom-permission \  
--aws-account-id AWSACCOUNTID
```

To test the custom permissions that are applied to a role or user, log in to the user's account. When a user logs into Quick, they are granted the highest privilege role that they have access to. The highest privileged role a user can be granted is Admin. The lowest privileged role that a user can be granted is Reader. For more information about roles in Quick, see [Managing user access inside Quick](#).

Verifying permissions for a user

After you assign a custom permissions profile, you can verify which profile is active for any user and at which level it applies.

To verify a user's active permissions profile

1. Open the [Quick console](#).
2. Choose **Manage Quick**.
3. In the left navigation, choose **Permissions**, and then choose **Custom permissions**.
4. Choose **Check permissions**.
5. Enter the username of the user you want to verify.
6. Quick displays the active profile name and the level at which it applies (user, role, or account).

Use this feature to confirm that Deny by Default profiles are correctly applied and that precedence is working as expected.

Deny by Default

Important

Test Deny by Default configurations in a development or staging account before applying them to production users. Verify that the behavior matches your expectations before broad deployment.

Overview

What is Deny by Default?

Without Deny by Default, any new capability that Amazon Quick ships is automatically available to all users upon release. Administrators must manually restrict each new capability after every release. This reactive approach can expose regulated environments to unapproved capabilities before administrators have time to evaluate them.

With Deny by Default enabled, new capabilities in a restricted category are automatically denied on launch day without administrator action. Administrators can then evaluate each new capability and explicitly allow it when ready. Restricting a category also restricts the capabilities that already exist in it, so you allow the ones your users need.

How category restriction works

When you restrict a capability category in a custom permissions profile, Quick restricts every capability in that category for assigned users. This includes capabilities that Quick launches in the future. To keep specific capabilities available, allow them explicitly in the profile.

Existing profiles are not changed

Restricting a category affects only the profile you configure it in. Existing custom permissions profiles do not inherit the restriction. To use it, either create a new profile with the category restricted, or edit an existing profile. If you enable it on an existing profile, all capabilities in that category become restricted – including capabilities the profile previously allowed. Allow the ones your users need.

When to use this feature

Deny by Default is designed for organizations that require explicit approval before new capabilities reach users. Common use cases include:

- Financial services organizations with model risk management (MRM) policies that require review of AI capabilities before adoption
- Healthcare organizations that must evaluate capabilities for compliance before making them available to users
- Enterprises running controlled rollouts where new features are introduced incrementally
- Organizations with regulatory or compliance requirements that mandate review before capability adoption

Key concepts

The following table describes the key terms and concepts for Deny by Default.

Term	Definition
Custom permissions profile	A named configuration object that defines which Amazon Quick capabilities are restricted for a set of users or roles. You can assign profiles at the user, role, or account level.

Term	Definition
Capability category	A named grouping of Amazon Quick capabilities used as the unit of control for Deny by Default. At launch, the supported category is AI (all AI and LLM-powered capabilities).
Deny by Default	A setting you can enable within a custom permissions profile to restrict all capabilities in a designated category. Capabilities that Quick launches in that category later are also restricted automatically, without requiring administrator action at launch time.
Category restriction	Restricting a capability category denies every capability in that category for users assigned to the profile, current and future, except capabilities explicitly allowed in the profile.
Profile scope	Restricting a category affects only the profile it is configured in. Existing custom permissions profiles do not inherit the restriction.
Precedence hierarchy	Evaluation order when a user has permissions at multiple levels: User overrides Role, which overrides Account. The most specific level takes precedence.
DefaultCategoryEffects	The API field that specifies Deny by Default behavior per capability category. <code>DENY_BY_DEFAULT</code> restricts all capabilities in the category—current and future—except those explicitly allowed. When omitted, the default behavior is allow by default.
Conflict resolution	When a capability is denied at one level but explicitly allowed at another, the most specific permission (User > Role > Account) wins.

Prerequisites

Before you configure Deny by Default, verify that the following prerequisites are met.

- **IAM permissions** – You must be a Quick administrator with the following IAM permissions:
 - `quicksight:CreateCustomPermissions`
 - `quicksight:UpdateCustomPermissions`
 - `quicksight:DescribeCustomPermissions`
 - `quicksight:ListCustomPermissions`
 - `quicksight>DeleteCustomPermissions`
 - `quicksight>DeleteAccountCustomPermission`
 - `quicksight>DeleteRoleCustomPermission`
 - `quicksight>DeleteUserCustomPermission`
 - `quicksight:DescribeAccountCustomPermission`
 - `quicksight:DescribeRoleCustomPermission`
 - `quicksight:ListCustomPermissionAssignments`
 - `quicksight:UpdateAccountCustomPermission`
 - `quicksight:UpdateRoleCustomPermission`
 - `quicksight:UpdateUserCustomPermission`
- **Identity configuration** – Your Quick account must be integrated with IAM Identity Center, Active Directory, or configured with Quick managed users.

Note

All API operations use Amazon QuickSight naming conventions. Permission strings use the `quicksight:` prefix.

How capabilities are evaluated

The following rules describe how capability evaluation works when Deny by Default is enabled. The rules are evaluated in order.

1. If a capability belongs to a restricted category (for example, AI) with `DENY_BY_DEFAULT` and is not listed in Capabilities, the user is denied by default. This is the primary effect of Deny by Default.

2. If a capability is explicitly listed in Capabilities with ALLOW, the user is permitted for that specific capability. This overrides the DENY_BY_DEFAULT category default for that capability only. Other capabilities in the same category remain denied unless they are also explicitly set to ALLOW.
3. If a capability is explicitly listed in Capabilities with DENY, the user is denied.
4. If a capability is not listed and does not belong to a restricted category, the user is permitted (allow-by-default behavior is unchanged).

Quick evaluates these rules each time a user accesses a capability, based on the capability's category membership at that time. This is why a restricted category also covers capabilities that are added later.

This means that when Quick launches a new capability in a restricted category, that capability is automatically denied for all users whose profile has DENY_BY_DEFAULT enabled for that category. No administrator action is required. The restriction takes effect on the day the capability launches.

 Note

Custom permission profiles are resolved in this priority order: user-level, then role-level, then account-level (first match wins). A user with a user-level allow-by-default profile is not subject to an account-level Deny by Default profile. Assign Deny by Default profiles at the appropriate level for your requirements.

Supported categories

The following table describes the supported categories for Deny by Default.

Category	Covers
AI	All AI and LLM-powered capabilities in Quick, including chat agents, flows, spaces, knowledge bases, apps AI inference, Q analyses, and Quick desktop AI features.

Note

The category system is modeled using tags. A capability can carry multiple category tags. If a capability is tagged with more than one category and more than one has Deny by Default enabled, the most restrictive setting applies.

Configuring Deny by Default (Quick console)

You configure Deny by Default as part of creating or editing a custom permissions profile. The setting is controlled by the **Restrict capabilities** section on the profile configuration page.

To enable Deny by Default on a custom permissions profile

1. Step 1: Open the custom permissions settings

- a. Open the [Quick console](#).
- b. Choose **Manage Quick**.
- c. In the left navigation, choose **Permissions**, and then choose **Custom permissions**.
- d. Create a new profile by choosing **Create profile**, or edit an existing profile by choosing the actions menu and then **Edit**.

2. Step 2: Configure the restriction

- a. In the **Restrict capabilities** section, turn on the toggle for the category you want to restrict (for example, **Restrict AI capabilities**). Turning on this toggle enables Deny by Default for that category. Any capability in this category that isn't explicitly allowed in the profile is automatically denied for users assigned to this profile.
- b. In the **Capabilities & features** section, explicitly allow any capabilities within the restricted category that you want users to retain access to. Only capabilities you explicitly allow are available. All others in the category, including any new capabilities that Quick launches in the future, are denied.
- c. Review the live preview panel on the right to verify your configuration matches your intent.

3. Step 3: Save the profile

Choose **Create** or **Update** to save the profile.

To remove Deny by Default from a profile, edit the profile and turn off the category toggle in the **Restrict capabilities** section. Capabilities that were previously denied become available to users assigned to the profile.

Configuring Deny by Default (AWS CLI)

You can use the AWS Command Line Interface to create and manage custom permissions profiles with Deny by Default enabled.

The following example creates a custom permissions profile with the AI category set to DENY_BY_DEFAULT, with the ChatAgent capability explicitly allowed.

```
aws quicksight create-custom-permissions \  
--aws-account-id AWSACCOUNTID \  
--custom-permissions-name PERMISSIONNAME \  
--capabilities '{"ChatAgent": "ALLOW"}' \  
--governance '{"DefaultCategoryEffects": {"AI": "DENY_BY_DEFAULT"}}'
```

The following example updates the Deny by Default settings on an existing profile.

```
aws quicksight update-custom-permissions \  
--aws-account-id AWSACCOUNTID \  
--custom-permissions-name PERMISSIONNAME \  
--capabilities '{"ChatAgent": "ALLOW", "Flow": "ALLOW"}' \  
--governance '{"DefaultCategoryEffects": {"AI": "DENY_BY_DEFAULT"}}'
```

Important

UpdateCustomPermissions performs a full replacement. You must send all --capabilities and --governance values, not just the delta. Any capability or setting you omit from the update call is removed from the profile.

The following example describes a profile with Deny by Default enabled.

```
aws quicksight describe-custom-permissions \  
--aws-account-id AWSACCOUNTID \  
--custom-permissions-name PERMISSIONNAME
```

Note

DescribeCustomPermissions returns the Governance field only if the profile was created or updated with Deny by Default settings. Profiles without Deny by Default return only the Capabilities field.

To remove Deny by Default from a profile entirely using the CLI, omit the `--governance` flag on the update call.

Examples**Example 1: Converting an existing profile to Deny by Default**

This example shows how to convert an existing allow-by-default profile that explicitly denies certain AI capabilities to a Deny by Default profile.

Scenario: You have a custom permissions profile that explicitly denies Flow, Automate, and ChatAgent capabilities. All other capabilities are available by default.

Run the following command to view your current profile configuration:

```
aws quicksight describe-custom-permissions \
--aws-account-id AWSACCOUNTID \
--custom-permissions-name "RestrictAI-Finance"
```

The command returns output similar to the following:

```
{
  "CustomPermissions": {
    "CustomPermissionsName": "RestrictAI-Finance",
    "Capabilities": {
      "Flow": "DENY",
      "Automate": "DENY",
      "ChatAgent": "DENY"
    }
  }
}
```

The problem: When Quick launches a new AI capability, it is automatically available to users with this profile because the profile only blocks explicitly listed capabilities. New capabilities are not listed, so they are allowed by default.

The solution: Convert the profile to Deny by Default. When you convert, you invert the logic. Instead of listing capabilities to deny, you list capabilities to allow. Everything not explicitly allowed is denied.

First, determine which AI capabilities you want to allow. In this case, you want users to retain access to Research, Topic, KnowledgeBase, and Space.

Run the following command to update the profile with Deny by Default:

```
aws quicksight update-custom-permissions \  
--aws-account-id AWSACCOUNTID \  
--custom-permissions-name "RestrictAI-Finance" \  
--capabilities '{"Research": "ALLOW", "Topic": "ALLOW", "KnowledgeBase": "ALLOW",  
"Space": "ALLOW"}' \  
--governance '{"DefaultCategoryEffects": {"AI": "DENY_BY_DEFAULT"}}'
```

Result: Now when Quick launches any new AI capability, it is automatically denied for users with this profile. Only Research, Topic, KnowledgeBase, and Space are available. Flow, Automate, ChatAgent, and any future AI capabilities are denied.

Important

Remember that UpdateCustomPermissions performs a full replacement. Include all capabilities you want to set, not just the ones related to the conversion.

Example 2: Setting up Deny by Default for a new account

This example shows how to lock down all AI capabilities from day one for a new Quick account.

Scenario: You are onboarding a new account and want to restrict all AI capabilities until your team evaluates each one individually.

Run the following command to create a profile that denies all AI capabilities by default with no explicit allows:

```
aws quicksight create-custom-permissions \  

```

```
--aws-account-id AWSACCOUNTID \  
--custom-permissions-name "DenyAllAI-NewAccount" \  
--governance '{"DefaultCategoryEffects": {"AI": "DENY_BY_DEFAULT"}}'
```

Run the following command to assign the profile at the account level so it applies to all users:

```
aws quicksight update-account-custom-permission \  
--aws-account-id AWSACCOUNTID \  
--custom-permissions-name "DenyAllAI-NewAccount"
```

Result: All AI capabilities are denied for every user in the account. As you evaluate and approve individual capabilities, you can update the profile to add them to the allow list.

Example 3: Allowing a specific new capability after launch

This example shows how to approve a new capability that Quick has launched after you enabled Deny by Default.

Scenario: Quick launches a new AI capability called NewAIFeature. Your team evaluates it and decides to approve it for users.

First, describe the current profile to get the current state:

```
aws quicksight describe-custom-permissions \  
--aws-account-id AWSACCOUNTID \  
--custom-permissions-name "RestrictAI-Finance"
```

The command returns the following output:

```
{  
  "CustomPermissions": {  
    "CustomPermissionsName": "RestrictAI-Finance",  
    "Capabilities": {  
      "Research": "ALLOW",  
      "Topic": "ALLOW",  
      "KnowledgeBase": "ALLOW",  
      "Space": "ALLOW"  
    },  
    "Governance": {  
      "DefaultCategoryEffects": {  
        "AI": "DENY_BY_DEFAULT"  
      }  
    }  
  }  
}
```

```
    }
  }
}
```

Update the profile to include the new capability in the allow list. You must include all existing capabilities plus the new one, because the update performs a full replacement:

```
aws quicksight update-custom-permissions \
--aws-account-id AWSACCOUNTID \
--custom-permissions-name "RestrictAI-Finance" \
--capabilities '{"Research": "ALLOW", "Topic": "ALLOW", "KnowledgeBase": "ALLOW",
"Space": "ALLOW", "NewAIFeature": "ALLOW"}' \
--governance '{"DefaultCategoryEffects": {"AI": "DENY_BY_DEFAULT"}}'
```

Result: The new capability is now available to users with this profile. All other unlisted AI capabilities remain denied.

Example 4: Multi-level precedence with Deny by Default

This example shows how Deny by Default interacts with the precedence hierarchy when profiles are assigned at multiple levels.

Scenario: You want most users locked down with Deny by Default for AI, but you want authors to have access to more capabilities. A specific data scientist needs full access to all AI capabilities.

Account-level (most restrictive): Create a profile that denies all AI by default with no allows. Assign it at the account level.

```
aws quicksight create-custom-permissions \
--aws-account-id AWSACCOUNTID \
--custom-permissions-name "Account-DenyAllAI" \
--governance '{"DefaultCategoryEffects": {"AI": "DENY_BY_DEFAULT"}}'

aws quicksight update-account-custom-permission \
--aws-account-id AWSACCOUNTID \
--custom-permissions-name "Account-DenyAllAI"
```

Role-level (Author): Create a profile with Deny by Default but with more capabilities allowed for power users. Assign it to the Author role.

```
aws quicksight create-custom-permissions \
```

```
--aws-account-id AWSACCOUNTID \  
--custom-permissions-name "Author-LimitedAI" \  
--capabilities '{"Research": "ALLOW", "Topic": "ALLOW", "ChatAgent": "ALLOW"}' \  
--governance '{"DefaultCategoryEffects": {"AI": "DENY_BY_DEFAULT"}}'\  
  
aws quicksight update-role-custom-permission \  
--role AUTHOR \  
--aws-account-id AWSACCOUNTID \  
--namespace default \  
--custom-permissions-name "Author-LimitedAI"
```

User-level (data scientist): Create an allow-by-default profile with no restrictions. Assign it to the specific user.

```
aws quicksight create-custom-permissions \  
--aws-account-id AWSACCOUNTID \  
--custom-permissions-name "DataScientist-FullAccess"\  
  
aws quicksight update-user-custom-permission \  
--aws-account-id AWSACCOUNTID \  
--namespace default \  
--user-name DATA_SCIENTIST_USERNAME \  
--custom-permissions-name "DataScientist-FullAccess"
```

Result:

- Most users (Readers, Admins without user-level overrides) receive the account-level profile with all AI capabilities denied.
- Authors receive the role-level profile with Research, Topic, and ChatAgent allowed. All other AI capabilities are denied.
- The data scientist receives the user-level profile with no restrictions. They are not subject to the account-level or role-level Deny by Default settings because user-level takes highest precedence.

Troubleshooting

A user can still access a capability I expected to be denied

Verify the following:

- Use the Check permissions feature to verify that the correct profile is assigned to the user.

- Check the precedence hierarchy. A user-level or role-level profile might override your account-level Deny by Default profile.
- Check if the capability is explicitly set to ALLOW in the profile.

Limit profiles

Intended audience: System administrators and Amazon Quick administrators

With Limits Management, you can monitor resource consumption. You can cap per-user usage before account-level service limits are reached. You can govern two resource types: *index storage* and *agent hours*.

The mechanism is straightforward: you create reusable *limit profiles* that define per-user caps, then assign those profiles to individual users, roles, or the entire account as a default. Quick enforces the most specific assignment that applies to each user.

Understanding index storage and agent hours

Limit profiles govern two resource types: index storage and agent hours. Understanding how each is measured helps you set effective per-user caps.

Index storage

Index capacity is measured by the original (raw) file size of documents at their source location before any processing by Quick. Supported source locations include Amazon S3, SharePoint, and Confluence. For full details on how index capacity is allocated, pooled, and purchased, see [Index capacity](#).

Index storage allocation is based on user count and subscription tier and is pooled at the payer-account level. Because the pool is shared, a single heavy user can consume capacity that would otherwise be available to other users. With limit profiles, you can cap individual consumption so that no single user exhausts the shared pool.

Agent hours

Agent hours measure the time consumed by AI-powered features in Quick. The following activities are metered:

- Quick Research – Time from research request submission until the report completes.
- Quick Flows – Runtime while workflows are actively executing, excluding idle and waiting time.
- Quick Automate – Runtime of deployed automations.
- Desktop app – Time spent actively using the AI assistant.
- Custom apps – Runtime of AI-powered custom applications.
- Artifact generation – Time spent generating documents, presentations, and images.

Viewing usage

Index capacity consumption is visible three ways: by user, by knowledge base, and by space. The **Per-user Capacity Overview** table provides an at-a-glance summary.

The **Per-user Capacity Overview** table provides an at-a-glance summary with the following columns.

User

The user's display name.

Role

The user's assigned role (Author or Reader).

Storage usage

Current index storage that the user consumed.

Agent-hours usage

Agent hours that the user consumed in the current billing cycle.

Usage distribution

A breakdown of consumption across knowledge bases and spaces.

User detail side panel

Select a user row to open a side panel that shows the following details:

- Linked limit profiles for each resource type
- Storage consumption compared to the user's effective limit

- Agent-hours consumption for the current billing cycle
- Usage distribution across knowledge bases and spaces

For additional analytics on agent-hours consumption trends, see [Usage metrics](#).

Managing limit profiles

Create a limit profile

Each limit profile requires a name and at least one limit. You can configure the following fields:

- **Name** (required) – A descriptive name for the profile.
- **Description** (optional) – Additional context about the profile's purpose.
- **Storage limit** (optional) – The maximum index storage in GB that the user can consume.
- **Agent hours per month** (optional) – The maximum agent hours the user can consume per billing cycle.

You must set at least one of the two limits (storage or agent hours) when creating a profile.

Assign a limit profile

You can assign a limit profile to:

- An individual user
- A role (Author or Reader)
- The account default (applies to all users who do not have a more specific assignment)

Resolution hierarchy

When multiple assignments apply to a user, the most specific assignment determines the effective limit, in the following priority order:

1. **User-level** – A limit profile assigned directly to a user takes highest priority.
2. **Role-level** – A limit profile assigned to the user's role (Author or Reader) applies if no user-level assignment exists.
3. **Account-level** – The account default profile applies if neither a user-level nor role-level assignment exists.

4. **System default** – The built-in subscription entitlement applies when no admin-configured profile is in effect.

Group-level assignment (assigning a profile to an identity-provider group) is planned for a future release. When available, group-level assignment will slot between user-level and role-level in the resolution hierarchy.

How limits are enforced

As long as a user stays below their effective limit, they can work without interruption.

When a user reaches 100% of a limit:

- **Index storage** – New file uploads and knowledge base ingestion are blocked for that user.
- **Agent hours** – New agent invocations are blocked until the next billing cycle begins.

Existing content is always preserved. Tightening a limit never deletes files or knowledge bases.

Enforcement is global across all AWS Regions. Usage in any Region counts toward the same per-user limit.

End-user visibility

Your users can view their own consumption in the **My usage** widget (available in **My stuff**). The widget shows:

- Agent hours used compared to their monthly limit. This value resets at the start of each billing cycle.
- Index storage capacity consumed compared to their effective limit.

For programmatic management of limit profiles, see the Amazon Quick API Reference.

Governance

With governance features in Quick, you control actions that affect shared assets. Use approval workflows to require review before the approver grants access.

Topics

- [Approval workflows](#)
- [Data loss prevention](#)

Approval workflows

Applies to: Enterprise Edition

Intended audience: System administrators and Amazon Quick administrators

Approval workflow overview

With approval workflows, you require a designated approver group to review and approve certain actions on assets in Amazon Quick before those actions take effect. Use approval workflows when you want a trusted reviewer to validate a sharing decision before the approver grants access. This reduces the risk of unauthorized access to sensitive assets.

Currently, approval workflows support sharing actions. The target user receives access only after the approver approves the request.

Approval workflows support the following asset types:

- Knowledge bases
- Spaces
- Custom chat agents

How it works

Approval workflows are opt-in. Sharing requires approval only after an administrator creates and enables a policy for an asset type.

An approval workflow involves three personas:

Administrator

Creates and manages approval policies.

Requester

Submits share requests that require approval.

Approver

Reviews and acts on pending approval requests.

For administrators: setting up approval policies

As an administrator, you create and manage approval policies in the **Approval Policies** section of the account management console.

To navigate to approval policies

1. Sign in to Quick as an administrator.
2. Choose your account name, and then choose **Manage Account**.
3. In the left navigation pane, choose **Approval Policies**.

To create an approval policy

1. Choose **Create Policy**.
2. Enter a policy name and an optional description.
3. Select the asset types that the policy applies to.
4. Choose one or more approver groups. These are existing identity groups from IAM Identity Center, IAM federation, or Active Directory.
5. In **Assign Policy**, select the user groups whose members must go through the approval workflow when sharing.
6. Choose **Create Policy**.

To manage existing policies, choose a policy to view its details in the side panel. From there, you can edit the policy name, description, asset types, and approver groups. You can also delete policies that are no longer needed.

For requesters: submitting a share request

Approval requirements depend on the sharing action:

Add a new person to an asset

Requires approval.

Upgrade access (for example, Viewer to Owner)

Requires approval.

Downgrade access

Takes effect immediately; no approval required.

Remove access

Takes effect immediately; no approval required.

To submit a share request

1. Locate the asset that you want to share.
2. Choose **Share**.
3. Add a person or group, or change an existing role.
4. If an approval policy is active, an approval request form appears with the following fields:
 - (Required) **Notes** – Context for the approver.
 - (Optional) **Severity** – For example, Low, Medium, or High.
 - (Required) **Approval needed by** – The date by which a decision is needed. This must be a future date.
5. Choose **Send Request**.

What happens after submission

After you submit a request, the following occurs:

- The approver group receives viewer-level access to the asset so they can review it. Quick grants this access when you submit the request and does not remove it automatically afterward. The asset owner removes the approver group manually.
- The request enters *Pending* status.
- The target user does not receive access until the request is approved.

Tracking requests

To track your submitted requests, navigate to **My stuff**, and then choose the **My Tasks** widget. Choose **Submitted By Me**.

Opening a request shows the status (Pending, Approved, or Denied), the asset and target user, your submitted notes, severity, date, and any approver comments or decision history.

For approvers: reviewing and acting on requests

To access pending requests, navigate to **My stuff**, and then choose **My Tasks**. You can view requests under **All** and **Assigned To Me** (requests routed to your approver group).

Before you can act on a request, you must claim it. Choose **Claim Request** to move the request to **Assigned To Me**. This exposes the **Approve** and **Deny** actions.

The access that approvers receive depends on the asset type:

Executable assets (agents)

Execution access to test. The asset runs in the creator's context, so the approver does not gain direct access to underlying data sources.

Content assets (spaces, knowledge bases)

Read access to review content.

After reviewing the asset, choose one of the following actions:

- **Approve** – The target user receives access. The request status changes to Closed – Approved.
- **Deny** – The request is declined. The requester receives reason codes and written feedback. The request status changes to Closed – Denied.

Verifying access after approval

After the approver approves a request, you can verify that Quick granted access correctly.

- **Target user** – The shared asset appears under the relevant section. Refresh the page if it does not appear immediately.
- **Requester** – In **My stuff**, choose **My Tasks**, and then choose **Submitted By Me**. The request shows *Approved* status. The details panel shows the history and timestamp.

- **Administrator** – Check the asset's sharing settings. The audit trail in AWS CloudTrail captures Submit, Approve, Deny, and Revoke actions with user, asset, timestamps, and notes.

Package sharing for custom chat agents

When you share a custom chat agent, you can submit a package share request that includes the agent and its dependencies (knowledge bases, connectors, and spaces) as a single all-or-nothing approval.

- **Approve** – The target user gets the agent and all listed dependencies.
- **Deny** – No access to any component is granted.

Approvers can view the full dependency list on the asset page before deciding.

API reference

For programmatic management of approval policies, see the Amazon Quick API Reference.

Frequently asked questions

Is approval required by default?

No. Approval workflows are opt-in only. No approval is required unless an administrator creates and enables a policy.

Which edition supports approval workflows?

Approval workflows are available in Quick Enterprise edition.

Who can be an approver?

Both Professional and Enterprise users can be approvers. Professional and Enterprise are user subscriptions within Quick Enterprise edition, not separate editions of Quick.

What groups can I use as approver groups?

You can use existing identity groups from IAM Identity Center, IAM federation, or Active Directory.

What happens if a request is denied?

The requester receives reason codes and written feedback. They can revise and resubmit the request.

Is approver access to the asset removed after a decision?

No. The approver group retains viewer-level access. The asset owner removes the approver group manually.

What happens if I delete an approval policy?

In-flight pending requests continue to completion. New share requests for those asset types no longer require approval.

Can a requester cancel a pending request?

Yes. In **My stuff**, choose **My Tasks**, choose **Submitted By Me**, and then choose **Cancel request**.

Are approvers notified of new requests?

Yes. Approvers are notified when new requests are submitted.

Can I audit approval events?

Yes. All approval events (Submit, Approve, Deny, and Revoke) are logged in AWS CloudTrail.

Data loss prevention

Intended audience: System administrators and Amazon Quick administrators

Amazon Quick integrates with Microsoft Purview to enforce your organization's data loss prevention (DLP) policies. After you connect your Microsoft Purview tenant, Quick classifies and protects sensitive files shared through spaces, chat, and knowledge bases. Sensitive data stays within the boundaries that your organization's information governance policies define.

Data loss prevention concepts

Before you configure DLP in Amazon Quick, familiarize yourself with the following concepts.

Sensitivity labels

Sensitivity labels are classifications that you define in your Microsoft Purview tenant to describe the sensitivity level of content – for example, *Public*, *Confidential*, or *Highly Confidential*. You can organize labels hierarchically with parent labels and sub-labels (child labels).

Enforcement actions

Enforcement actions define how Quick handles a file that carries a specific sensitivity label.

- **Block** – Quick rejects the file, and the file can't be shared or uploaded. You receive an error message.
- **Warn** – In interactive capabilities, you see a warning message and can choose to proceed or cancel the upload.
- **Allow** – Quick shares the file without restriction.

Default action

The default action is the enforcement action that Quick applies to any file whose sensitivity label you haven't explicitly mapped to a specific action. You set the default action during configuration. The default action also covers new labels that you publish in Microsoft Purview after setup. Quick governs files that carry those labels right away, so nothing goes unprotected. To give a new label its own action, edit the configuration and map it explicitly.

Provider outage action

The provider outage action determines how Quick handles files when the DLP provider (Microsoft Purview) is temporarily unavailable. You can set it to one of the following actions.

- **Block** – Quick rejects files until the provider is available again.
- **Warn** – Quick permits files through and shows a warning during the outage.
- **Allow** – Quick permits files through without restriction.

Capabilities

Capabilities are the areas within Amazon Quick where Quick applies DLP enforcement.

- **Spaces** – Files that users upload to collaborative workspaces.
- **Chat** – Files that users share through chat conversations.
- **Knowledge bases** – Content synced through SharePoint, OneDrive, and other connectors.

Label inheritance

In the label mapping configuration, you can assign a parent label an explicit action (Block or Warn) or leave it set to the default action. Sub-labels (child labels) inherit their parent label's action. A label that you haven't explicitly mapped follows its mapped parent label's action if one exists. Otherwise, your default action applies.

Prerequisites

Prepare the following before you create a DLP configuration.

Microsoft Purview sensitivity labels

You must already publish your sensitivity labels (for example, *Public*, *Confidential*, and *Highly Confidential*) in your Microsoft Purview or Microsoft 365 tenant. Quick reads these labels live from your tenant.

A Microsoft Entra ID app registration

Register an application in Microsoft Entra ID (Azure AD) that Quick uses to read labels and classify files. Note the **Directory (tenant) ID** and **Application (client) ID**. Both are GUIDs that you need to create a client secret credential. You must grant the app registration the following API permissions.

API permission	Type	Description
UnifiedPolicy.Tenant.Read	Required	Read the unified data-protection policy for the tenant.
SensitivityLabels.Read.All	Required	Read all sensitivity labels defined in the tenant.
Files.Read.All	Optional	Required only if you use knowledge bases backed by SharePoint or OneDrive.

An AWS Secrets Manager secret

Store the Purview credentials in an AWS Secrets Manager secret in the same AWS account as your Quick account. Use the following JSON format for the secret value.

```
{
  "clientId": "abcdef12-3456-7890-abcd-ef1234567890",
  "clientSecret": "<client secret value>",
  "tenantId": "11111111-1111-1111-1111-111111111111"
}
```

Note

Quick stores only the ARN of your secret. Quick reads your credentials at runtime and never displays them in the console, returns them from the API, or writes them to logs.

⚠ Important

Use exactly these three key names: `clientId`, `clientSecret`, and `tenantId`. Quick looks up these keys verbatim, so a renamed or missing key causes validation to fail with a missing-fields error. Store the value as a single JSON object.

Access for Quick to your secret

Sign in to your AWS account as an administrator and open the Quick admin console. Then grant Quick permission to read your Secrets Manager secret.

1. Choose **Permissions**, and then choose **AWS resources**.
2. Select **AWS Secrets Manager** from the list of services.
3. Choose your secret.
4. Choose **Save**.

Secrets encrypted with a customer managed KMS key

If a customer managed KMS key encrypts your Secrets Manager secret, Quick also needs `kms:Decrypt` permission on that key, in addition to `secretsmanager:GetSecretValue`.

⚠ Important

The **Permissions, AWS resources** page can't grant `kms:Decrypt`. Edit the Quick service role directly in IAM to add it. After you edit the role manually, you can no longer manage it from the **AWS resources** page in the admin console.

Add a statement like the following to the role. Replace the account ID, secret name, and key ID with your own.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [ "secretsmanager:GetSecretValue" ],
```

```
"Resource": [
  "arn:aws:secretsmanager:<home-region>:<account-id>:secret:<secret-name>",
  "arn:aws:secretsmanager:<other-region>:<account-id>:secret:<secret-name>"
],
{
  "Effect": "Allow",
  "Action": [ "kms:Decrypt" ],
  "Resource": [ "arn:aws:kms:<other-region>:<account-id>:key/<key-id>" ]
}
```

Replicating your secret across Regions

Quick reads your secret in every AWS Region where it evaluates files, so the secret must exist in each of those Regions. Replicate it by using **Replicate secret to other Regions** in the Secrets Manager console, and make sure the role grants access to every replica ARN.

If a customer managed KMS key doesn't encrypt the secret, replicating it is sufficient. If a customer managed KMS key encrypts the secret, use one of the following approaches.

Option 1: Single-Region keys

1. Replicate the secret to the other Region.
2. Detach the replica from the primary secret so that it becomes standalone.
3. Create a KMS key in the other Region.
4. Attach that key to the standalone secret.
5. Grant `kms:Decrypt` on the new key, as described in [the section called "Secrets encrypted with a customer managed KMS key"](#).

Option 2: Multi-Region key

1. Create a multi-Region key in your home Region.
2. Set the multi-Region key as the encryption key on the secret.
3. Replicate the multi-Region key to the other Region.
4. Replicate the secret to that Region, and select the replicated multi-Region key as its encryption key.

5. Grant kms:Decrypt on the replica key, as described in [the section called "Secrets encrypted with a customer managed KMS key"](#).

Creating a DLP configuration

Sign in to the admin console. Choose **Governance**, choose **Data loss prevention**, and then choose **Create DLP configuration** to open the setup wizard. The wizard has four steps, listed in the navigation pane. The following table describes each step.

Step	Name	Description
1	Provider	Select the DLP provider (Microsoft Purview).
2	Credentials	Name the configuration and connect to AWS Secrets Manager.
3	Label mapping	Scope enforcement and map sensitivity labels to actions.
4	Review	Confirm settings and activate the configuration.

Step 1: Provider

Select **Microsoft Purview** as the DLP provider, and then choose **Continue**.

Step 2: Credentials

Enter the following information:

- **Configuration name** – A descriptive name that the dashboard shows – for example, Knowledge base DLP.
- **Secret ARN** – Paste the ARN of your Secrets Manager secret.
- Choose **Validate** to connect to Purview with the secret and check the required read permissions.

 Note

Each required permission shows a green check when validation succeeds. If a permission is missing or the secret is invalid, Quick flags the issue inline so that you can correct it and retry.

 Important

You can't continue to the next step until validation succeeds. Verify that the app registration has all required API permissions and that the secret ARN is correct before you retry.

Step 3: Label mapping

Map each sensitivity label to an action:

- **Default action** – Choose the action that Quick applies to any label that you don't explicitly map. The options are **Block**, **Warn**, or **Allow**. A conservative starting point is **Block**.
- **Sensitivity labels** – Quick loads the current list of labels from your Purview tenant. For each label, choose its action. Sub-labels inherit their parent's action. Choose **Refresh** if you recently changed labels in Purview and want to pull the latest list.
- **Provider outage action** – Choose what Quick does when Purview is unreachable: **Block** (fail closed – safest), **Warn**, or **Allow** (fail open). For a knowledge base that holds sensitive material, **Block** is the recommended provider outage action.

Choose **Review**.

The following table describes the enforcement actions.

Allow

The file proceeds normally. Quick applies no restrictions.

Warn

The file proceeds, but Quick notifies you in interactive capabilities (Chat and Spaces) if it carries a sensitive label.

Block

Quick rejects the file and doesn't ingest it into the Quick index.

Step 4: Review

Review the summary – the provider, configuration name, authentication method, capabilities, the label-to-action breakdown, the default action, and the provider outage action. Choose the edit link next to any section to return to it.

Use the **Active / Inactive** toggle to set whether Quick enforces the configuration as soon as you create it (the default is **Active**). Quick saves an **Inactive** configuration but doesn't enforce it until you activate it. When you're ready, choose **Create**.

Your configuration now appears as a card on the **Data loss prevention** dashboard. If you leave it **Active**, Quick begins enforcing on new file uploads to your knowledge bases.

How enforcement works

When a file enters a Quick capability that an active DLP configuration protects, the following sequence occurs.

1. A file enters a protected Quick capability – for example, Quick ingests documents during a knowledge base sync.
2. Quick DLP retrieves the file's Microsoft Purview sensitivity label.
3. Quick matches the label against the label-to-action mapping in your configuration.
4. Quick applies the resulting action.
 - **Allow** – The file proceeds normally.
 - **Warn** – The file proceeds, but Quick notifies you in interactive capabilities (Chat and Spaces) if it carries a sensitive label.
 - **Block** – Quick rejects the file.
5. For a label that you haven't explicitly mapped, Quick checks for a mapped parent label. If a mapped parent exists, Quick uses that parent's action. Otherwise, it applies your **default action**. New labels that you publish in Purview after setup are also covered by the default action, so Quick governs them immediately. To give a new label its own action, edit the configuration and map it explicitly.
6. If Quick can't reach Purview, it applies your configured **provider outage action**.

 Tip

For a knowledge base that holds sensitive material, we recommend that you set both the default action and the provider outage action to **Block**. This prevents Quick from ingesting unlabeled files or files that arrive during a Purview outage.

Enforcement by capability

DLP enforcement applies to three capabilities: **Chat**, **Spaces**, and **Knowledge bases**. The same actions apply in all three. What you see, and where Quick records the result, depends on whether Quick scans the file interactively or during a background sync.

Interactive capabilities: Chat and Spaces

Quick evaluates files as you share or upload them, while you wait for the result.

- **Block** – Quick rejects the file, and you receive an error message.
- **Warn** – You see a warning message and can choose to proceed or cancel.
- **Allow** – Quick shares the file without restriction.

Knowledge bases

Quick evaluates files during ingestion, after a sync starts, so no user waits for the outcome. Two consequences follow.

- **There is no warn prompt.** Quick ingests a file that carries a label mapped to **Warn**, and the **Warn** action has no user-visible effect on this capability.
- **Results appear in the Quick Observability report.** To see what DLP did during a knowledge base sync, open the Quick Observability report for that sync. The report shows one status per file.

BLOCKED

Quick didn't ingest the file into the index. A **BLOCKED** status can indicate DLP enforcement or an internal failure. The error message for internal failures is generic.

ADDED

Quick ingested a new file. This status includes files that match **Warn** labels; the report doesn't show them separately.

MODIFIED

Quick updated and re-ingested an existing file. Files that match **Warn** labels appear as **MODIFIED** if the file changed since the last sync.

Testing your configuration

From the configuration's card on the dashboard, choose the **Test** (file-scan) icon. Provide the Amazon S3 URI (for example, `s3://bucket/key`) of a labeled file, and confirm that it returns the expected Allow, Warn, or Block result before you rely on it in production.

Note

The test action performs a dry run using the current label-to-action mapping. It doesn't modify the file or affect any live data.

Managing configurations

From each dashboard card, you can perform the following actions.

- **Edit** (pencil icon) – Change credentials, capabilities, label mappings, the default action, or the provider outage action.
- **Activate / Deactivate** (toggle) – Turn enforcement on or off. Deactivating asks for confirmation, because it stops protection for the covered capabilities.
- **Delete** – Remove the configuration entirely (available from the edit view).

Note

The dashboard card warns you when your Purview tenant has new labels that you haven't mapped yet. Until you map them, they follow your default action.

Limits and validations

The following limits apply to DLP in Amazon Quick.

- **DLP configurations per account** – 10.
- **Active DLP configurations per account** – 1.
- **Maximum file size scanned** – 500 MB. Quick doesn't scan larger files.

Supported file types

Quick scans the file types that Microsoft Purview supports for labeling: Word, Excel, PowerPoint, PDF, Visio, Project, common image formats (JPEG, PNG, TIFF, DNG, PSD), XPS, Power BI (.pbix), .dwfx, and email (.msg, .eml). For the authoritative list, see [supported file types](#) on the Microsoft website.

One distinction matters when you read this list: only file types that Microsoft Purview supports *for labeling* are in scope. Quick doesn't scan files that require Azure Rights Management to decrypt, or file types that apply only to protected labels.

Protected labels

If a file carries a label that has protection enabled, Quick doesn't scan the file and applies your default action instead. Protected labels don't appear in the label-mapping step in the admin console. You can still supply them through the `CreateDlpSetting` API, but the evaluation outcome is the same: the default action applies.

When the default action applies

Beyond labels that you haven't explicitly mapped, Quick falls back to your default action when any of the following is true:

- The file carries no sensitivity label.
- The file type isn't supported (see [the section called "Supported file types"](#)).
- The file exceeds the 500 MB size limit.
- The file is empty (0 bytes).
- The file carries a protected label.

Note

In the public API, the default action is the `unmappedAction` field.

Security best practices

- **Use a dedicated app registration for DLP.** Register a separate Microsoft Entra ID application for Quick DLP rather than reusing an app registration that serves other integrations. A dedicated registration lets you rotate or revoke its credentials independently and keeps its permission scope narrow.
- **Grant only the permissions that DLP needs.** Scope the app registration to `SensitivityLabels.Read.All` and `UnifiedPolicy.Tenant.Read.AddFiles.Read.All` only if you use knowledge bases backed by SharePoint or OneDrive; it isn't required to list or map labels.
- **Keep credentials in Secrets Manager.** Quick stores only the ARN of your secret. Quick reads your credentials at runtime and never displays them in the console, returns them from the API, or writes them to logs.
- **Monitor data loss prevention (DLP) activity.** Quick delivers a `DLP_LOGS` vended log to CloudWatch Logs, Amazon S3, or Firehose. Use it to audit DLP policy changes, including when enforcement was enabled or disabled and which principal made the change (`DLP_SETTING_*` events, status, and `last_updated_by`). You can also identify blocked, warned, or inspection-failed files by filtering on `event_type` and `failure_type`. For setup and the full schema, see [the section called "Monitoring using CloudWatch Logs"](#).

API reference

For programmatic management of DLP settings, see the Amazon Quick API Reference.

Setting up Quick knowledge bases

To set up Quick ingestion knowledge bases, follow the instructions for each connector:

- [the section called "Knowledge base integration"](#)
- [the section called "Knowledge base integration"](#)
- [the section called "Knowledge base integration"](#)
- [the section called "Web Crawler integration"](#)

- [the section called "S3 integration"](#)
- [the section called "Knowledge base integration"](#)

Troubleshooting

The following table lists common issues and how to resolve them.

Validation fails with a permission error

Verify that the app registration has `UnifiedPolicy.Tenant.Read` and `SensitivityLabels.Read.All` granted and admin-consented.

Secret ARN not found

Confirm that the Secrets Manager secret exists in the same AWS account as Quick, and that you granted Quick access under **Permissions, AWS resources**.

The labels list is empty or outdated

Use **Refresh** in Step 3 (Label mapping) to re-fetch the current label list from your Purview tenant.

Files are unexpectedly blocked during a knowledge base sync

Check the Quick Observability report for the sync. A **BLOCKED** status can also indicate an internal failure rather than DLP enforcement. Check whether the file carries a label that you mapped to Block, whether the default action is Block, and the provider outage action if Purview was unreachable during the sync.

Customizations

Using Amazon Quick, you can create a customized experience for people using either the AWS Management Console or Quick consoles embedded in your application.

Currently, different options for customizing Quick are available separately in the console and the Quick API. Following, you can find information about the available options.

The following customization options are currently available:

- You can customize the welcome content Quick provides for new users:
 - You can accept or decline the sample assets. These assets include sample datasets and analyses that are added when a person signs in for the first time.

- You can show or hide default introductory videos. These videos include the animation that displays for new users and also the tutorial videos shown on the Amazon Quick home page.
- You can create and specify a default theme.
- You can customize dashboard report emails, pixel perfect report emails, and alert emails by editing the email template.

Important

All customizations apply only to the AWS Region that you are using in the API or that is selected in the Amazon Quick console.

To check your Region setting, you can use one of the following procedures.

To check your AWS Region on the Amazon Quick console

1. Choose your profile icon to open the menu.
2. View your current AWS Region, listed next to a location icon.
3. (Optional) Choose another AWS Region from the menu to change to that Region. Remember to change back after you are finished with customizations.

To check your AWS Region using the AWS CLI

- On the command line, enter the following command and press Enter to view the current settings.

```
aws configure list
```

To reconfigure your default Region, use the `aws configure` command.

To keep your default Region, you can add the `--region` parameter to most CLI commands.

Topics

- [Amazon Quick brand customization](#)
- [Chat agent customization in Amazon Quick](#)

Amazon Quick brand customization

Amazon Quick allows account admins to customize their application's branding and visual theme to align with their organization's guidelines. This customization includes the following visual elements to create a cohesive look and feel across all non-administrative Amazon Quick console pages, schedules, alerts, and email reports.

- Logo
- Favicon
- Associated alt text for visual assets

The following list shows the different areas customizable theme colors are grouped into.

Brand colors

- **Global navigation bar** colors are applied to the topmost bar in the Amazon Quick UI and include the company logo that is displayed in the standard and embedded Amazon Quick consoles.
- **Application bar** colors are applied to the secondary navigation bar that contains contextual actions.

Interaction colors

- **Accent** colors are applied to interactive elements like buttons, borders, and icons.

Surface colors

- **Primary** colors are applied to high-emphasis surfaces like the homepage background and text.
- **Secondary** colors are applied to practical surfaces like borders, backgrounds, and form fields. Secondary colors are used alongside primary colors.

Status colors

- **Success** colors are applied to success messages.
- **Danger** colors are applied to error messages.
- **Warning** colors are applied to warning messages.
- **Info** colors are applied to informational messages.

Data visualization colors

- **Dimension** colors are used to identify associations between data columns that share the same role.
- **Measure** colors are used to identify metrics or measured values.

Features

- **Visualization** colors are applied to the **Visualization** icon.
- **Insight** colors are applied to the **Insight** icon.
- **Connection** colors are applied to the **Connection** icon.
- **Automation** colors are applied to the **Automation** icon.

Use the following sections to get started with brand customization in Amazon Quick.

Topics

- [Permissions for Amazon Quick brand customization](#)
- [Create a custom brand in Amazon Quick](#)
- [Setting a default theme for Amazon Quick analyses with the Amazon Quick APIs](#)

Permissions for Amazon Quick brand customization

To set up a brand, you must be granted an Admin role through IAM Identity Center or IAM. Admins whose roles are granted to them within Amazon Quick can't create brands. To learn more about integrating your account with IAM Identity Center, see [Configure your Amazon Quick account with IAM Identity Center](#). For information about admin roles and capabilities, see [Understanding Amazon Quick subscriptions and roles](#).

Admin users can only manage brands that are in the same capacity Region as their Amazon Quick account.

The IAM role that you use to create a brand in Amazon Quick must contain `quicksight:*` or granular action permissions to manage brands in the admin console. The following granular permissions are required for admins to work with Amazon Quick brands:

- `quicksight:CreateBrand`

- quicksight:UpdateBrand
- quicksight:DescribeBrand
- quicksight:DescribeBrandPublishedVersion
- quicksight:UpdateBrandPublishedVersion
- quicksight>DeleteBrand
- quicksight:ListBrands
- quicksight:UpdateBrandAssignment
- quicksight:DescribeBrandAssignment
- quicksight>DeleteBrandAssignment

After you confirm that your Admin role contains the required permissions, you can [Create a custom brand](#) in the Amazon Quick admin console.

Create a custom brand in Amazon Quick

Use the following procedure to create a custom brand in Amazon Quick.


1. Open the [Quick console](#).
2. Choose the user icon at the top right, and then choose **Manage Quick**.
3. Choose **Customize application**.
4. On the **Customize application** page that opens, choose **ADD BRAND**. The **Brand settings** page opens.
5. Navigate to the **Brand Info** section.
6. For **Brand name**, enter a name for the brand. The brand name can contain up to 512 characters.
7. (Optional) For **Brand description**, enter a description for the custom brand. The brand description can contain up to 512 characters.
8. Navigate to the **Logo** section.

Logo

Primary



Favicon



Alt text

9. For **Primary**, choose the ellipsis (three dots) next to the primary icon, and then choose **Replace image**.
10. In the **Choose image** pop up that opens, perform one of the following actions:
 - a. Drag and drop image into the **Drag an image here** box.
 - b. Choose **Select a file** to select a file from your computer.
 - c. Enter a public URL or Amazon S3 URI in the text bar.

The image that you choose must be a .jpeg, .png, or .svg format and can't exceed 1MB.

When you are finished choosing an image, choose **Apply**.

11. For **Favicon**, choose the ellipsis (three dots) next to the favicon, and then choose **Replace image**.
12. In the **Choose image** pop up that opens, perform one of the following actions:
 - a. Drag and drop image into the **Drag an image here** box.
 - b. Choose **Select a file** to select a file from your computer.
 - c. Enter a public URL or Amazon S3 URI in the text bar.

The image that you choose must be a .jpeg, .png, or .svg format and can't exceed 1MB.

When you are finished choosing an image, choose **Apply**.

13. (Optional) For **Alt text**, enter alt text to display with the logo. The alt text can contain up to 512 characters.

14. To make changes to the theme colors of the brand, navigate to the **Appearance** pane on the left and choose **Theme**.
15. The **Theme settings** page appears and displays all parts of a Amazon Quick theme that can be customized. The following image shows the configuration settings of the global navigation bar.



16. To change the background color of an area, navigate to the item that you want to change and choose the **Background** color swatch.
17. In the **Custom color** pop up that appears, choose a color from the color gradient or enter a hex code value in the **HEX** bar, and then choose **APPLY**.
18. To change the foreground color of an area, navigate to the item that you want to change and choose the **Foreground** color swatch.
19. In the **Custom color** pop up that appears, choose a color from the color gradient or enter a hex code value in the **HEX** bar, and then choose **APPLY**.
20. When you are finished configuring a custom brand, choose **PUBLISH** to publish and apply the brand customization to all Amazon Quick user accounts. If you don't want to publish the brand, choose **SAVE** to save the brand for later.

When you finish creating a brand in Amazon Quick, the new brand appears in the brands table on the **Customize application** page of the Quick admin console. The **Status** column of the brands table indicates which brand is currently published to the Quick account. To make changes to a custom brand, locate the brand that you want to change in the brands table, choose the ellipsis (three dots) icon in the **Actions** column, and then choose **Publish**, **Edit**, or **Delete**.

Once you publish a brand, it can take up to 10 minutes for the new brand to propagate across all users.

Setting a default theme for Amazon Quick analyses with the Amazon Quick APIs

To set a default theme by using the API

1. Identify the custom theme that you want to use as the default, and locate its theme ID. If you want to use one of the QuickSight starter themes, skip this step.

To get the theme ID of a custom theme, use the [ListThemes](#) API operation for the Region where the theme is. Make sure that the theme is in the same Region with the users or groups that need to use it.

The following example shows a shell script that uses the `list-themes` command in the AWS CLI. It sets the AWS account ID and the AWS Region as variables. If you previously used `aws configure` to set a default Region, adding the `--region` variable to your command overrides your default setting.

```
#declare variables
awsacct1='111122223333'
region='us-west-2'

aws quicksight list-themes \
--region $region \
--aws-account-id $awsacct1 \
--type 'CUSTOM'
```

2. Use the [ListUsers](#) or [ListGroups](#) API operation to collect the Amazon Resource Names (ARNs) for users or groups that need to use the theme as a default. You need only the top-level ARN. If all your users are part of the same group, use the group ARN.

For more information on Amazon Quick ARNs, see [ARN formats](#) in the *Quick API Reference*.

3. If you're using a custom theme, grant access to the theme for the ARNs that you collected in the previous step. If you're using a starter theme, skip this step because all users have access to starter themes.

The following example shows a shell script that uses the [update-theme-permissions](#) command. The `grant-permissions` parameter is shown using shorthand syntax. You can use JSON or YAML instead. For more information, see [Specifying parameter values](#) in the *AWS Command Line Interface User Guide*.

```
#declare variables
awsacct1='111122223333'
namespace='default'
region='us-west-2'
theme-id='bdb844d0-0fe9-4d9d-b520-0fe602d93639' #Find this with list-themes

aws quicksight update-theme-permissions \
```

```
#Specify region if necessary: --region $region \
--aws-account-id $awsacct1 \
--theme-id $theme-id \
--grant-permissions
Principal="arn:aws:quicksight:$region:$awsacct1:group/$namespace/
QuickSight_Group_Name",Actions="quicksight:DescribeTheme","quicksight:ListThemeVersions","
```

4. Assign the theme as the default for the same ARN or ARNs.

```
#declare variables
awsacct1='111122223333'
namespace='default'
region='us-west-2'
theme-id='bdb844d0-0fe9-4d9d-b520-0fe602d93639'

aws quicksight create-account-customization \
#Specify region if necessary: --region $region \
--aws-account-id $awsacct1 \
--namespace $namespace \
--account-customization
DefaultTheme="arn:aws:quicksight:$region:$awsacct1:theme/$theme-id"
```

Currently, there are three starter themes: Classic, Midnight, and Seaside. Their ARNs are the capitalized spelling of their theme name. If you are using a starter theme instead of a custom theme, use one of the following theme ARNs:

- arn:aws:quicksight::aws:theme/CLASSIC
- arn:aws:quicksight::aws:theme/MIDNIGHT
- arn:aws:quicksight::aws:theme/SEASIDE
- arn:aws:quicksight::aws:theme/RAINIER

Chat agent customization in Amazon Quick

When a user account subscribes to Quick, the service automatically creates a system default chat agent. This agent powers the default chat experience, allowing Amazon Quick users to leverage all chat functionalities (such as uploading files to chat, using large language model (LLM) parametric knowledge, and answering from their enterprise data) out of the box. The system default chat agent can be updated in the Amazon Quick agent section by a select set of users who are designated as owners for the agent by the admin.

Amazon Quick uses multiple safety controls for chat interactions in the Amazon Quick web experience. You can configure blocked words and phrases under **Guardrails and safety controls**. For information about built-in safeguards and blocked words and phrases, see [AI guardrails in Amazon Quick](#).

Admins can also configure whether URLs in chat responses appear as clickable hyperlinks or plain text. This setting applies across all chat agents and flows in your instance, allowing you to control how links are presented to users.

Note

Admins can also control permissions for whether users can create and use chat agents and flows. For instructions on how to do that, see [Custom permissions](#).

The following sections outline how to edit the system default agent and configure clickable external links in chat responses.

Topics

- [Grant user permissions to edit system default chat agent](#)
- [Edit system chat agent settings](#)
- [Configure clickable external links in chat responses](#)

Grant user permissions to edit system default chat agent

The system default chat agent can be edited by users the admin designates as owners to this agent in the admin console. The following procedure shows you how to grant admin permissions to a user so that they can edit the system default agent.

To grant permissions to edit system default agent

1. Log in to the Amazon Quick console.
2. Select **Manage Quick** from the Amazon Quick admin console.
3. From the admin console left navigation menu, select **Account**, and then select **Manage assets**.
4. In **Manage assets**, select **Chat agents**.
5. Select **My Assistant** from the list of chat agents shown. Then, from the **Actions** menu, select **Share**.

6. In the **Share assets** modal, use the **User or Group** search bar to find the user or group you want to designate as owners for the system agent.
7. Then, from **Permissions**, choose **Owner**. Then, select **Share**.

Edit system chat agent settings

To customize your system chat agent, users designated as owners need to login to Amazon Quick and visit the agent library page. Admins can directly access this page by clicking on the **Go to agent** link in admin console under **Chat agent customization**. Once in Amazon Quick, follow these steps to edit your system default chat agent's properties. You can preview and test how your system agent works as you configure it.

To update system chat agent settings

1. Log in to the Amazon Quick console.
 2. Select **Manage Quick** from the Amazon Quick admin console.
 3. From the admin console left navigation menu, select **Customization**, and then select **Chat agent customization**.
 4. In **Chat agent customization**, in **System chat agent**, select **Go to chat agent**.
 5. In the **My Assistant** page, in **Configure chat agent**, customize the following sections:
 - a. In **Chat agent persona**, configure your chat agent's personality, identity, tone, and response style. For detailed information about agent customization options, see [Working with chat agents](#) in the *Amazon Quick User Guide*.
-  **Note**
- This agent powers default chat interactions across all users of this account with chat feature access. Ensure that your instructions (e.g. identity and response style) work well for all users of this account.
- b. In **Reference documents**, upload files that remain active in your chat agent's memory to guide all interactions. For detailed information about file upload options, see [Working with chat agents](#) in the *Amazon Quick User Guide*.

Note

Since this agent is meant for broad use, this is a place to upload enterprise-level response templates and guides.

The system default agent isn't linked to specific knowledge sources, actions, or spaces to ensure it works for all users regardless of their access permissions. This setting can't be changed.

6. In **Customization**, configure details to help recognize the chat agent. For detailed information about customization options, see [Working with chat agents](#) in the *Amazon Quick User Guide*.

Configure clickable external links in chat responses

Admins can configure whether URLs in chat agent responses appear as clickable hyperlinks. The following procedure shows you how to enable clickable hyperlinks for all chat agents in your Amazon Quick instance.

To enable clickable external links

1. Log in to the Amazon Quick console.
2. Select **Manage Quick** from the Amazon Quick admin console.
3. From the admin console left navigation menu, select **Customization**, and then select **Chat agent customization**.
4. Under **Clickable external links**, turn the toggle on.

Using the Amazon Quick analytics dashboard

Analytics in Amazon Quick provides insights into how users interact with your instance across all features and capabilities. The analytics dashboard delivers the data-driven insights you need to make informed decisions about security, performance, and user experience.

Note

The analytics dashboard is only available for IAM administrators.

The Amazon Quick analytics system captures detailed usage metrics from chat conversations, space interactions, flow executions, action invocations, research activities, automations, and custom agent usage. This comprehensive data collection enables you to understand not just what users are doing, but how effectively they're leveraging Quick to accomplish their goals while maintaining security and compliance standards.

You can use analytics to:

- Monitor user adoption and engagement across all Quick features
- Identify potential issues and monitor system health through available analytics
- Analyze user feedback and satisfaction trends to improve the user experience
- Understand resource utilization patterns and identify your most valuable content
- Make data-driven decisions about feature deployment, training needs, and resource allocation
- Monitor security-related events such as blocked messages and failed queries

Topics

- [Analytics access and permissions](#)
- [Analytics dashboard overview](#)
- [Accessing the analytics dashboard](#)
- [Understanding analytics metrics](#)
- [Analytics limitations](#)

Analytics access and permissions

Access to analytics in Amazon Quick is controlled by user roles and permissions to ensure data security and privacy. Currently, only IAM administrators have access to the analytics dashboard.

To access the analytics dashboard, you need the following permission:
quicksight:QuickSuiteUsageMetrics

User type	Analytics capabilities
IAM administrators	• View all analytics data for the entire instance • Access detailed query and response data • Monitor team adoption and engagement metrics

Important

If you have analytics access, you can view detailed query text, failed-query content, and the full conversation context associated with negative feedback. Grant analytics access only to authorized administrators, and apply your organization's data-handling and incident-response policies to analytics review workflows.

Analytics dashboard overview

The Amazon Quick analytics dashboard is organized into several key sections that provide different perspectives on your instance usage and security events:

Dashboard filters and controls

The analytics dashboard provides filtering options to focus on specific data:

- **Agent filter** – Select from two available options: "All Agents" to see aggregate data across all agents, or "My Assistant (System Agent)" to view system agent metrics
- **Date range picker** – Choose different time periods to analyze trends and patterns

Active user metrics

Track how users interact with your Amazon Quick instance over time.

This section displays user activity metrics including:

- **Active users** – Daily, weekly, and monthly active user counts with trend comparisons
- **Usage trends** – Growth patterns and engagement changes over time
- **Feature adoption** – Usage distribution across chat, spaces, flows, and other features

Use these metrics to understand user adoption patterns and identify engagement opportunities.

Chat agent and asset metrics

Monitor chat usage patterns and user feedback to understand engagement with conversational features and identify areas for improvement.

Conversation analytics

This section provides insights into chat usage including:

- **Total conversations** – Number of chat sessions initiated
- **Total queries** – Number of messages sent by users
- **Average queries per user** – Engagement depth metrics
- **Average queries per conversation** – Conversation length patterns displayed in a time-series chart showing trends over months
- **Conversations, queries, and active users chart** – Interactive line chart displaying total conversations, total queries, and active users trends over time with monthly granularity

These metrics help optimize chat experiences and identify usage trends.

User feedback analytics

Track user satisfaction to continuously improve system quality.

This section displays user feedback data including:

- **Positive feedback** – Thumbs up interactions and trends
- **Negative feedback analysis** – Thumbs down interactions with detailed reasons and breakdown by category
- **Customer feedback trend chart** – Time-series visualization showing positive vs. negative feedback patterns over months
- **Query scope selection chart** – Analytics showing how users select different query scopes, with breakdowns by "All Resources Selected", "No Resources Selected", and "Specific Resources"
- **Feedback trends** – Satisfaction changes over time with 3-month historical data
- **Thumbs down details** – View specific queries that received negative feedback, including the full conversation context and user-provided reasons
- **Thumbs down reason breakdown** – Detailed categorization of negative feedback reasons to identify common issues and improvement areas

Leverage feedback data to enhance user satisfaction and system performance.

Query issue analysis

Identify and resolve query problems to improve system performance.

The query issue analysis section helps you identify and resolve problems:

- **Issue percentage breakdown** – View the percentage of queries that failed by issue type
- **Failed query details** – See the actual query text from users, ordered by most recent, with issue type classification
- **Unanswered queries section** – Dedicated analytics for queries that received no response, helping identify system gaps and user needs
- **Full conversation context** – Click on any failed query to see the complete conversation leading up to the issue
- **Issue trends** – Track how query issues change over time to identify patterns

Use this analysis to proactively address system issues and improve user experience.

Asset creation and usage

Monitor resource utilization to optimize content and feature adoption.

This section provides insights into resource utilization including:

- **Assets created** – New spaces, agents, flows, actions, research, and automations
- **Assets used** – Active resource engagement
- **Number of users by asset type chart** – Horizontal bar chart showing user distribution across different asset types including Actions, Automations, Chat, Custom Agents, Flows, Research, and Spaces
- **Most used assets** – Most popular content and tools (spaces, agents, actions, flows, research)

Use these insights to guide content strategy and resource allocation decisions.

Automations, flows, and research agent hours usage

Monitor agent hours usage to manage subscription limits and optimize resource allocation across different subscription tiers.

This section includes the following tabs:

- **Automation and Flows Agent hours** – Provides insights into agent hours consumed by automations and flows
- **Research Agent hours** – Provides insights into agent hours consumed by research activities

Agent hours consumption metrics include:

- **Total Professional hours** – Agent hours consumed by Professional tier users
- **Total Enterprise hours** – Agent hours consumed by Enterprise tier users
- **Additional hours** – Overage hours used beyond included monthly allocations
- **Hours used by subscription type chart** – Horizontal bar chart showing user distribution by percentage of included hours consumed (25-49%, 50-74%, 75-99%, >100%) for both Enterprise (4hrs/mo) and Professional (2hrs/mo) tiers
- **Additional agent hours used by week** – Weekly trend chart showing additional hours consumed by Enterprise users after exceeding their included hours

Use these metrics to understand user adoption patterns and identify engagement opportunities.

Accessing the analytics dashboard

You can access the analytics dashboard from the Amazon Quick console if you have the appropriate permissions.

To access the analytics dashboard

1. Log in to the Amazon Quick console.
2. From the left navigation menu, select **Analytics**, and then select **Amazon Quick Usage Metrics**.

The analytics dashboard opens, displaying an overview of your instance usage and available analytics.

3. Use the date picker to specify the time interval for the graphs. You can choose from the following preset intervals:
 - **Last week** – This interval starts from the previous Sunday at 00:00 and goes until the following Sunday at 00:00 in the UTC time zone.
 - **This week** – This interval starts on Monday at 00:00 and goes until 00:00 today as per the timezone relevant for the primary region of the Amazon Quick account.

- **Last month** – This interval begins on the 1st day of the last month at 00:00 and ends on the first day of this month at 00:00 in the UTC time zone.
- **This month** – This interval begins on the 1st day of this month at 00:00 and goes until 00:00 today in the UTC time zone.

Note

If you don't see the Analytics option in the navigation menu, contact your instance administrator to request access.

Understanding analytics metrics

The analytics dashboard displays various metrics that help you understand usage patterns, performance, and security events. The following sections explain key metrics in detail.

User activity metrics

User activity metrics help you understand how many people are actively using your Amazon Quick instance:

- **Daily active users** – Number of unique users who sent at least one message in a given day
- **Weekly active users** – Number of unique users who sent at least one message in the current week starting from Monday
- **Monthly active users** – Number of unique users who sent at least one message in the current month starting from the 1st day of the month

Note

Active user metrics are based on chat interactions and may not reflect all types of Quick usage.

Security monitoring metrics

Security monitoring metrics help you track potential security issues and policy enforcement:

- **Blocked messages** – Messages blocked by guardrails for safety, policy, or content violations

- **Failed messages** – Messages that encountered errors, with specific error codes for troubleshooting

Analytics limitations

The following limitations apply to Amazon Quick analytics:

- Analytics primarily covers chat interactions and may not capture all types of Quick usage
- Data processing may have a delay of up to 24 hours
- Historical data availability depends on when analytics was first enabled
- Some metrics may be aggregated to protect individual user privacy
- Analytics data retention follows your instance's data retention policies
- Detailed query and response data is only available to instance administrators

Note

For comprehensive monitoring beyond chat interactions, consider using CloudWatch metrics and logs in conjunction with the analytics dashboard. For more information, see [Monitoring Quick](#).

AWS security in Quick

Amazon Quick provides a secure platform that enables you to distribute dashboards and insights to tens of thousands of users, with multiple-region availability and built-in redundancy.

Cloud security at AWS is the highest priority. As an AWS customer, you benefit from a data center and network architecture that is built to meet the requirements of the most security-sensitive organizations.

Security is a shared responsibility between AWS and you. The [shared responsibility model](#) describes this as security *of* the cloud and security *in* the cloud:

- **Security of the cloud** – AWS is responsible for protecting the infrastructure that runs AWS services in the AWS Cloud. AWS also provides you with services that you can use securely. The effectiveness of our security is regularly tested and verified by third-party auditors as part of the

[AWS compliance programs](#). To learn about the compliance programs that apply to Quick, see [AWS Services in Scope by Compliance Program](#).

- **Security in the cloud** – Your responsibility is determined by the AWS service that you use. You are also responsible for other factors, including the sensitivity of your data, your organization's requirements, and applicable laws and regulations.

This documentation helps you understand how to apply the shared responsibility model when using Amazon Quick. The following topics show you how to configure Amazon Quick to meet your security and compliance objectives. You also learn how to use other AWS services that can help you to monitor and secure your Amazon Quick resources.

Amazon Quick enables you to manage your users and content using a comprehensive set of security features. These include role-based access control, Microsoft Active Directory integration, AWS CloudTrail auditing, single sign-on using AWS Identity and Access Management (IAM) and third-party solutions, private VPC subnets, and data backup. Amazon Quick can also support FedRAMP, HIPAA, PCI DSS, ISO, and SOC compliance to help you meet industry-specific or regulatory requirements.

Topics

- [AI and agent security in Amazon Quick](#)
- [Data protection in Amazon Quick](#)
- [Incident response, logging, and monitoring in Amazon Quick](#)
- [Compliance validation for Amazon Quick](#)
- [Resilience in Amazon Quick](#)
- [Infrastructure security in Amazon Quick](#)
- [Security in Quick](#)
- [AWS managed policies for Amazon Quick](#)

AI and agent security in Amazon Quick

Amazon Quick includes AI capabilities that read content, take actions, and run automated workflows. These capabilities introduce properties that differ from traditional BI features. Use this page to understand what makes these capabilities different, how to apply the available controls, and where those controls have limits. For the corresponding security checklist, see [Security in Quick](#).

What makes these capabilities different

Three properties distinguish the AI and agent capabilities from traditional features such as dashboards and datasets.

- **The assistant reads content your users did not author** – Knowledge bases, web search, third-party research providers, and connectors all supply content that originates outside your organization. Any content from outside your trust boundary is a potential source of adversarial input. For the full list of content sources, see [Two-condition model for adversarial content risk](#).
- **The assistant can take actions with effects outside the service** – Through connectors, Model Context Protocol (MCP) server tools, and web search queries, the assistant can cause changes in external systems. For the full inventory of outbound paths, see [Outbound connections and data egress in Amazon Quick](#).
- **Some capabilities run without a person watching** – Automated workflows execute on a schedule or in response to events, without an interactive user session. This changes how you evaluate the risk of each action. For more information, see [Autonomy and identity](#).

None of these properties is unique to Amazon Quick. They are common across AI-powered services. Together, they mean that the same security review you apply to a traditional dashboard does not fully address the risk profile of an agent or automated workflow.

Two-condition model for adversarial content risk

The risk of acting on adversarial content requires two conditions to be present at the same time:

1. The assistant can read content from outside your trust boundary.
2. The assistant can take a consequential action.

Narrowing either condition reduces your exposure. A trust boundary is the set of content sources that you consider safe and under your control. Any content that enters the assistant's context from outside that boundary is untrusted by definition.

Read surface

The following sources can supply content to the assistant's context. Each is a point where content you did not author can enter.

- Knowledge base and index retrieval (including bring-your-own-index from Amazon Q Business)

- Files uploaded to spaces
- Files uploaded in chat
- Web search results
- Third-party research provider data
- Content returned by connectors
- Content returned by remote MCP server tools

Retained context also contributes to the read surface. Past conversations and, on Quick Desktop, entities extracted from connected data sources persist across sessions and can inform later responses. For more information, see [Memory and retained context](#).

Not all of these sources carry equal risk. Files uploaded by your own users in chat are within your trust boundary. Content returned by a web search query to the open internet is not. Evaluate each source against your trust boundary when you review your configuration.

Action surface

The following paths allow the assistant to cause effects outside Amazon Quick. Each is a point where a consequential action can occur.

- On-demand actions through third-party and AWS service connectors
- Automated workflow actions through service-level connectors
- Remote MCP server tool invocations
- Web search queries (send context externally)
- Quick Apps bridge API calls (sandboxed)
- Code execution on Quick Desktop (file effects on the user's machine, within configured folder permissions; other code execution surfaces are sandboxed without external reach)
- Scheduled agent execution

For the complete inventory of outbound paths, including what data can leave through each one and who controls it, see [Outbound connections and data egress in Amazon Quick](#).

Autonomy and identity

Amazon Quick supports two invocation models with different security characteristics. The same connector or action carries different risk depending on which model invokes it.

On-demand actions

On-demand actions are user-initiated and interactive. They authenticate with the individual user's personal credentials and act as the invoking user. Each on-demand action is bounded by that user's connector permissions, and a person is present in the session when it executes.

Automated workflows

Automated workflows are scheduled or event-triggered and non-interactive. They authenticate at service level rather than as an individual user. Because they use service-level authentication, they are not bounded by an individual user's connector permissions and they run unattended.

The consequence for your review is that an automated workflow can reach any resource the service-level credential permits, regardless of which user created the workflow. No interactive session is present to observe what happens during execution. An action can carry limited risk when you invoke it on demand. The same action can carry broader risk when an automated workflow invokes it with service-level reach.

Adding a review step

To add a review step to an automated workflow, use the human-in-the-loop task center. For workflows where you want a person to approve an action before it executes, the task center pauses execution and collects approval before the workflow continues. For more information, see [Human-in-the-loop task center](#).

Human oversight

You can apply human approval gates at five points in the lifecycle of AI and agent capabilities. Together, they form a layered model for human oversight.

1. **Capability approval (Deny by Default)** – New AI capabilities in a restricted category are denied on launch until an administrator explicitly allows them. This gate applies before a capability is available to any user. For more information, see [Deny by Default](#).
2. **Build-time consent (Apps integration consent)** – When the Apps agent adds an integration to an app (connector, space, dashboard visual, or AI inference), it prompts the author for approval. The published app never includes unapproved integrations, and app viewers inherit the approved integration scope rather than broader access. For more information, see [Integration consent model](#).
3. **Share approval (Flows)** – Custom permissions control which users can review and approve flow sharing requests, and whether creators can share flows without approval. Administrators

configure these capabilities through custom permissions profiles. For more information, see [Custom permissions](#).

4. **Runtime review (human-in-the-loop task center)** – For workflows where you want a person to approve an action before it executes, the human-in-the-loop task center pauses execution and collects approval. This gate is optional and per-workflow. For more information, see [Human-in-the-loop task center](#).
5. **Per-action user control (Desktop, Preview)** – On Quick Desktop, each system tool can be enabled or disabled individually and supports a three-tier permission model – **Full Access**, **Read Only**, and **Ask Each Time** – with granularity for each operation. **Full Access** runs operations without prompting for confirmation. **Ask Each Time** confirms every operation. This places an approval decision in your hands at the moment of each action, for the tools you have enabled. For more information, see [Access levels](#) and [Managing permissions](#).

Where a per-action prompt applies, it is a user-configurable choice. On Desktop, you set a tool to **Ask Each Time** to confirm every operation, or to **Full Access** to run without prompting. In the web experience, on-demand actions execute when you trigger them, and automated workflows are non-interactive by design – the gates at capability, build, and share time do the work there.

Per-action confirmation is therefore available but not universal. It depends on the surface and on how you configure each tool.

For an automated workflow running unattended, no interactive session exists to confirm anything. For those workflows, the review effort belongs at approval time, when you evaluate what the automation can reach and what it can do.

Controls and limits for AI and agent capabilities

Amazon Quick provides the following controls relevant to AI and agent capabilities. Each control is described together with its limits.

Retrieval permission enforcement

For knowledge bases with document-level access control list (ACL) support enabled, Amazon Quick enforces synchronized document permissions on every retrieval. Supported integrations include Amazon S3, Google Drive, Microsoft SharePoint, Atlassian Confluence Cloud, and Microsoft OneDrive. If Amazon Quick cannot evaluate document permissions for a query, it returns no documents rather than unfiltered results.

Limit – ACL enforcement applies only to knowledge bases created with ACL support. Knowledge bases created without ACL support do not enforce document-level permissions. ACL configuration is permanent and cannot be changed after creation. For more information, see [Best practices for managing ACLs in knowledge bases](#).

Custom permissions and deny by default

Custom permissions restrict the capabilities available to each role. They cannot grant capabilities that a user does not already have. The most specific level wins: user-level overrides role-level, which overrides account-level. The deny-by-default option for the AI capability category restricts all AI capabilities, current and future, so that new capabilities are denied on launch day until you explicitly allow them.

Limit – Custom permissions require an Enterprise subscription. They scope down access but cannot extend it. For more information, see [Custom permissions](#) and [Deny by Default](#).

Organization-level controls and their limits

Custom permissions apply at the account level, and an administrator of that account can change them. If you manage many accounts with AWS Organizations, you can add organization-level controls that individual account administrators cannot override. These controls act on AWS API operations and account sign-up. Understand what they reach and what they do not.

- **Restrict who signs up for Quick** – A service control policy (SCP) controls which accounts, editions, and identity types can create a Quick subscription, so new accounts cannot start a subscription outside your approved set. For more information, see [Using service control policies to restrict Amazon Quick sign-up options](#).
- **Restrict the management API operations** – An SCP denies the Quick API operations that create and manage agents, flows, spaces, and knowledge bases. Where the policy applies, principals cannot create those resources through the AWS Command Line Interface, the AWS SDKs, or CloudFormation.
- **Restrict AI feature availability with custom permissions, not SCPs** – SCPs act on AWS API operations and account sign-up. They do not control which AI features a user sees or can use within the product. To restrict AI features for your users, use custom permissions with the deny-by-default option for the AI capability category. For more information, see [Deny by Default](#).

Limit – Each control enforces a different layer. An SCP controls account sign-up and API-based management across your organization, and no account administrator can override it. Custom

permissions control which AI features your users can use, and an account administrator can change them. To detect a custom permissions change, record the Quick management API operations with AWS CloudTrail and review them. For more information, see [Incident response, logging, and monitoring in Amazon Quick](#).

Tool and connector supply chain

Tools reach the AI assistant only through integrations that an administrator creates. MCP integrations require an Enterprise subscription. The integration model restricts the transport and discovery path for external tools.

- Only remote MCP servers are supported. Local studio connections are not supported.
- During connection, Amazon Quick discovers and registers the tools that the MCP server provides. After discovery completes, each tool is listed as an action that you can review and turn on. Tools do not become available automatically.
- You can connect to private MCP servers through a VPC connection that has network access to the server.
- Administrators can restrict individual connectors through custom permissions. Each connector supports three permissions: **Create and Update action**, **Share action**, and **Use action**. For more information, see [Custom permissions](#).

For more information about configuring an MCP integration, see [Model Context Protocol \(MCP\) integration](#).

Limit – Authentication options vary by connector type. For remote MCP servers, agent-to-agent connectors, and generic HTTP connectors, connecting without authentication is a selectable option. The integrity of a connection depends on how it is configured. The service does not evaluate the trustworthiness of a third-party MCP server or the behavior of the tools it exposes. That assessment is yours to make.

Guardrail screening and blocked words

Built-in content safety screening checks user requests for harmful content, including hate speech, insults, sexual content, violence, and misconduct. These checks also assess user requests for prompt attacks. You can configure up to 50 blocked words and phrases that Amazon Quick checks against both user requests and generated responses.

Limit – These controls help reduce risk. However, they do not replace your responsibility to configure access, permissions, agents, flows, and data sources for your use case. Blocked-word

checks use exact-match string comparison, not semantic analysis, and are limited to 50 entries of up to 36 characters each. Blocked-word checks apply only to chat agents and flows in the Amazon Quick web experience. For more information, see [AI guardrails in Amazon Quick](#).

Code execution isolation

The AI assistant can write and run Python code on four surfaces. Each surface runs code in a sandboxed environment designed to contain the code's effects. The surfaces differ in where code runs and what it can reach.

- **Quick Desktop** – Code runs in an OS-level sandbox (macOS Seatbelt or Windows AppContainer) on the user's machine. File access inside the sandbox follows the folder permissions the user configures, and a persistent namespace carries variables across code runs within a conversation. Pre-installed Python packages only; package installation (`pip install`) is not supported. For more information, see [Code execution](#).
- **Web experience chat** – Code runs in an isolated, service-managed sandbox with no internet access. Outbound network requests from the sandbox fail. The sandbox does not persist beyond the conversation.
- **Automated workflows (Quick Automate)** – Code actions and custom agent code run within a restricted Python environment. Code actions are not the path for reaching external systems: external API calls belong to REST API integrations, and file system or database access belongs to the appropriate integrations. For more information, see [the section called "Code"](#).
- **Office extensions** – The assistant creates and executes code within the Office application sandbox (Outlook, Excel, Word, and PowerPoint). For more information, see [Connect Amazon Quick to your existing tools with extensions](#).

The Execute code locally capability in custom permissions restricts code execution in Quick Desktop for the users a profile applies to. Extension capabilities, including code execution within Office applications, can be restricted through extension access controls. For more information, see [Custom permissions](#) and [Extension access](#).

Limit – On Desktop, the sandbox always has access to the system temporary directories regardless of the configured folder permissions. On Desktop, the persistent namespace means data loaded by an earlier code run remains available to later code runs in the same conversation. Sandboxing contains what code can reach; it does not evaluate what the code does with content already in the conversation context. Review generated code output like any other assistant output.

Logging and audit coverage

AWS CloudTrail records API-level activity (who did what). Vended logs record conversational content (what was said and what was responded). Together, they provide after-the-fact visibility into AI and agent activity.

Limit – Logging does not prevent any action. It provides visibility after the fact. The service does not include real-time alerting; you must configure CloudWatch alarms separately if you need near-real-time notification. Vended chat logs do not capture the content of source documents that retrieval returns. For more information, see [Incident response, logging, and monitoring in Amazon Quick](#).

Limits stated plainly

The following limitations apply to the AI and agent capabilities in Amazon Quick:

- **Screening reduces risk but does not replace configuration** – Built-in screening and guardrails help reduce risk. They do not replace your responsibility to configure access, permissions, agents, flows, and data sources for your use case.
- **No universal human review requirement** – Human review is not a system-wide requirement for actions. Per-action confirmation is available where a tool is configured for it (on Desktop, through the **Ask Each Time** access level), and the human-in-the-loop task center is available for workflows where you choose to add a review step. However, on-demand actions in the web experience execute immediately, and automated workflows execute without user intervention.
- **Authentication options vary by connector type** – The authentication methods available depend on the connector type. For remote MCP servers, agent-to-agent connectors, and generic HTTP connectors, connecting without authentication is a selectable option.

Memory and retained context

Chat memory personalizes responses by remembering your preferences, context, and past interactions. Chat agents reference stored information to provide more relevant responses without you repeating information in every conversation. For more information, see [Memory and response personalization](#).

On Quick Desktop, a My Context surface provides a Knowledge graph tab that shows entities and relationships extracted from connected data sources, and a Memory tab that shows facts, procedures, and patterns learned from conversations. For more information, see [My context](#).

Why retained context matters for security

Retained context is part of the read surface described in [Two-condition model for adversarial content risk](#). Unlike a single conversation, retained context persists across sessions. Information that the assistant remembers from a past interaction can influence responses and actions in a future session.

Controls

- **Isolation** – Memories are specific to the individual user and are not shared with other users.
- **User control** – You can review your chat memory through **Conversation History**. On Desktop, you can additionally manage whether Quick learns from your conversations, searches your conversation history, and extracts entities from connected services. You can also view, edit, and delete individual memories. For more information, see [Privacy controls](#).
- **Full reset (Desktop)** – The **Clear all data** option irreversibly removes all conversations, knowledge graph data, saved credentials, and user preferences. For more information, see [Clearing all data](#).
- **Administrator control** – Custom permissions include capabilities that restrict building and managing engrams (profiles of a user's writing style and patterns) and creating and accessing knowledge memory in Quick Desktop. For more information, see [Custom permissions](#).
- **No training** – Quick does not use your conversations, files, or personal context to train or improve AI models.

Limitations

- Chat memory is not supported for accounts configured with AWS KMS customer managed keys. An organization that requires customer managed keys does not get chat memory. For more information, see [Memory and response personalization](#).
- Quick Desktop is in Preview. The My Context surface, the Knowledge graph tab, the ability to view, edit, and delete individual memories, and the **Clear all data** option are Desktop features. In the web experience, you review chat memory through **Conversation History**. Confirm which controls are available on the surface you deploy rather than assuming parity between Desktop and web.

Where to go next

- [Security in Quick](#) – The security checklist with actionable practices for configuring Amazon Quick.

- [Outbound connections and data egress in Amazon Quick](#) – The complete inventory of outbound paths, what data can leave through each, and who controls it.
- [AI guardrails in Amazon Quick](#) – Built-in content safety screening and administrator-configured blocked words.
- [Custom permissions](#) – Role-based capability restriction and deny-by-default governance.
- [Best practices for managing ACLs in knowledge bases](#) – Document-level access control for knowledge bases, including fail-closed retrieval behavior.

Data protection in Amazon Quick

The AWS [shared responsibility model](#) applies to data protection in Amazon Quick. As described in this model, AWS is responsible for protecting the global infrastructure that runs all of the AWS Cloud. You are responsible for maintaining control over your content that is hosted on this infrastructure. You are also responsible for the security configuration and management tasks for the AWS services that you use. For more information about data privacy, see [Data Privacy FAQ](#). For information about data protection in Europe, see the [General Data Protection Regulation \(GDPR\) Center](#).

For data protection purposes, we recommend that you protect AWS account credentials and set up individual users with AWS IAM Identity Center or AWS Identity and Access Management (IAM). That way, each user is given only the permissions necessary to fulfill their job duties. We also recommend that you secure your data in the following ways:

- Use multi-factor authentication (MFA) with each account.
- Use SSL/TLS to communicate with AWS resources. We require TLS 1.2 and recommend TLS 1.3.
- Set up API and user activity logging with AWS CloudTrail. For information about using CloudTrail trails to capture AWS activities, see [Working with CloudTrail trails](#) in the *AWS CloudTrail User Guide*.
- Use AWS encryption solutions, along with all default security controls within AWS services.
- Use advanced managed security services such as Amazon Macie, which assists in discovering and securing sensitive data that is stored in Amazon S3.
- If you require FIPS 140-3 validated cryptographic modules when accessing AWS through a command line interface or an API, use a FIPS endpoint. For more information about the available FIPS endpoints, see [Federal Information Processing Standard \(FIPS\) 140-3](#).

We strongly recommend that you never put confidential or sensitive information, such as your customers' email addresses, into tags or free-form text fields such as a **Name** field. This includes when you work with Amazon Quick or other AWS services using the console, API, AWS CLI, or AWS SDKs. Any data that you enter into tags or free-form text fields used for names may be used for billing or diagnostic logs. If you provide a URL to an external server, we strongly recommend that you do not include credentials information in the URL to validate your request to that server.

Amazon Quick does not use customer data for training or improving underlying LLMs.

For information about safety controls for AI-powered chat interactions, see [AI guardrails in Amazon Quick](#).

Topics

- [Data encryption in Amazon Quick](#)
- [Inter-network traffic privacy in Amazon Quick](#)
- [Outbound connections and data egress in Amazon Quick](#)
- [User lifecycle and data handling in Amazon Quick](#)

Data encryption in Amazon Quick

Amazon Quick uses the following data encryption features:

- Encryption at rest
- Encryption in transit
- Key management

You can find more details about data encryption at rest and data encryption in transit in the following topics. For more information about key management in Amazon Quick, see [Encrypting your Amazon Quick data with AWS Key Management Service customer managed keys](#).

Topics

- [Encryption at rest](#)
- [Encryption in transit](#)

Encryption at rest

Amazon Quick securely stores your Amazon Quick metadata. This includes the following:

- Amazon Quick user data, including Amazon Quick user names, email addresses, and passwords. Amazon Quick administrators can view user names and emails, but each user's password is completely private to each user.
- Minimal data necessary to coordinate user identification with your Microsoft Active Directory or identity federation implementation (Federated Single Sign-On (IAM Identity Center) through Security Assertion Markup Language 2.0 (SAML 2.0)).
- Data source connection data.
- Amazon Quick encrypts your data source credentials (username and password) and OAuth tokens with your default customer managed KMS key when one is registered with Amazon Quick. Amazon Quick re-encrypts credentials and tokens with the current default customer managed KMS key each time you update the credential or refresh an OAuth token. If you do not register a default customer managed KMS key, Amazon Quick encrypts this information with a service-managed key.
- Names of your uploaded files, data source names, and data set names.
- Statistics that Amazon Quick uses to populate machine learning (ML) insights.
- Data indexed to support Amazon Q in Quick. This includes the following:
 - Topics
 - Metadata related to your dashboards
 - Your first index capacity purchase
 - Your first chat
 - Your first space creation
 - Your first knowledge base creation

Important

Quick protects data indexed to support Amazon Q with the Amazon Q data key. That key is set the first time Amazon Q data is created in your account and cannot be changed afterward. To use a customer managed KMS key for Amazon Q data, register it as your account default key before creating any Amazon Q data. For more information, see [Amazon Q data key](#) and [Customer managed KMS key scope](#).

Amazon Quick securely stores your Amazon Quick data. This includes the following:

- When no customer managed KMS key applies, data at rest in SPICE is encrypted using hardware block-level encryption with keys that AWS owns and manages. For information about using a customer managed KMS key with SPICE datasets, see [Customer managed KMS key scope](#).
- When no customer managed KMS key applies, other data at rest is encrypted with a service-managed key. This includes the following:
 - Sample values for filters.
 - User feedback on chat responses, including optional free-text comments. User feedback is not encrypted with your customer managed KMS key or the Amazon Q data key.

Report artifacts (including email reports) can be encrypted with your default customer managed KMS key when one is registered. For more information, see [Customer managed KMS key scope](#).

- Amazon Quick encrypts cached query results at rest. Amazon Quick caches query results from a SPICE dataset that your customer managed KMS key protects under that same key, and encrypts other cached results with a service-managed key.

When you delete a user, Amazon Quick permanently deletes that user's metadata. What happens to the assets that the user owned, and to data such as the user's conversation history, depends on how the user is removed. For more information, see [User lifecycle and data handling in Amazon Quick](#).

When you unsubscribe from Amazon Quick, all of your Amazon Quick data is completely and permanently deleted. This includes all metadata, any data you have in SPICE, and Amazon Q data, including the Quick index. This deletion is immediate; there is no retention period during which the data can be recovered.

Encryption in transit

Amazon Quick supports encryption for all data transfers. This includes transfers from the data source to SPICE, or from SPICE to the user interface. However, encryption isn't mandatory. For some databases, you can choose whether transfers from the data source are encrypted or not. Amazon Quick secures all encrypted transfers by using Secure Sockets Layer (SSL).

Inter-network traffic privacy in Amazon Quick

To use Amazon Quick, users need access to the internet. They also need access to a compatible browser or a mobile device with the Amazon Quick mobile app installed. They don't need access to the data sources they want to analyze. This access is handled inside Amazon Quick. User

connections to Amazon Quick are protected through the use of SSL. So that users can access Amazon Quick, allow access to HTTPS and Web Sockets Secure (wss://) protocol.

You can use a Microsoft AD connector and single sign-on (IAM Identity Center) in a corporate network environment. You can further restrict access through the identity provider. Optionally, you can also use MFA.

Amazon Quick accesses data sources by using connection information supplied by the data source owner in Amazon Quick. Connections are protected both between Amazon Quick and on-premises applications and between Amazon Quick and other AWS resources within the same AWS Region. For connections to any source, the data source must allow connections from Amazon Quick.

Traffic between service and on-premises clients and applications

You have two connectivity options between your private network and AWS:

- An AWS Site-to-Site VPN connection. For more information, see [What is AWS site-to-site VPN?](#)
- An Direct Connect connection. For more information, see [What is AWS direct connect?](#)

If you are using AWS API operations to interact with Amazon Quick through the network, clients must support Transport Layer Security (TLS) 1.2 or later. Clients must also support cipher suites with Perfect Forward Secrecy (PFS), such as Ephemeral Diffie-Hellman (DHE) or Elliptic Curve Diffie-Hellman Ephemeral (ECDHE). Most modern systems such as Java 7 and later support these modes. You must sign requests using an access key ID and a secret access key that are associated with an IAM principal, or you can use the [AWS Security Token Service \(STS\)](#) to generate temporary security credentials to sign requests.

Traffic between AWS resources in the same region

An Amazon Virtual Private Cloud (Amazon VPC) endpoint for Amazon Quick is a logical entity within a VPC that allows connectivity only to Amazon Quick. The VPC routes requests to Amazon Quick and routes responses back to the VPC. For more information, see the following:

- [VPC endpoints](#) in the *Amazon VPC User Guide*
- [Connecting to a Amazon VPC with Amazon Quick](#)

Outbound connections and data egress in Amazon Quick

Stored customer data in Amazon Quick does not leave its AWS Region. The documented exceptions are AI inference processed within a defined geography (cross-Region inference) and the outbound paths that you or your administrator enables. Every path described on this page is either off by default or requires explicit configuration before data can leave the service boundary.

For related guidance on evaluating these paths before you enable them, see [Security in Quick](#).

Code execution is not an outbound path in the web experience. Python code that the AI assistant writes and runs in web chat executes in an isolated sandbox with no internet access. Automation code actions run in a restricted Python environment. External API calls belong to REST API integrations, which are connectors described on this page. For more information, see [Code execution isolation](#).

The following table summarizes each outbound path, the data that can leave, and the control that enables it.

Outbound path	Data that can leave	What enables it
Connectors	Conversational context, user requests, and structured parameters sent to the connected third-party or AWS service	Administrator creates and configures the connector, selects authentication method and scopes; individual connectors can be restricted through custom permissions
Remote MCP servers	Tool invocation parameters derived from conversational context	Administrator creates the MCP integration (Enterprise subscription required)
Web search	Queries derived from user requests sent to external web search providers	Administrator enables web search for the organization
Third-party research providers	Research queries derived from the research objective	Administrator sets up each integration individually and shares it with authorized users; premium providers require a

Outbound path	Data that can leave	What enables it
		provisioned license; the Research custom permission removes all providers for a profile
Quick Apps (public access)	App content rendered to anonymous internet viewers	App owner enables public sharing (Free and Plus accounts only)
Embedding	Dashboard visuals, analytics, and generative BI content rendered in an external web application	Administrator or developer configures embedding with IAM API calls and domain allowlisting (Enterprise required)
Exports and downloads	Dashboard data (PDF, CSV), research reports, analysis data	User-initiated; administrator can restrict through custom permissions
Log and metric delivery	Chat conversations, user feedback, agent hours, index storage usage, knowledge base file sync results	Administrator configures delivery destinations (CloudWatch Logs, Amazon S3, Firehose)
Cross-Region inference	Inference inputs and outputs move between Regions within a defined geography	Automatic platform behavior; no outbound destination for you to review or restrict

Connectors

Connectors create secure connections between Amazon Quick and third-party or AWS services. After an administrator configures a connector, it sends conversational context, user requests, and structured parameters to the target service.

- **What triggers it** – On-demand connectors execute when a user initiates them. Automated workflow connectors execute on a schedule or in response to an event, without user interaction.
- **Who controls it** – An administrator creates and configures each connector, choosing the authentication method (managed OAuth, custom user OAuth, API key, or service-to-service) and the scopes granted. A user can invoke only the connectors that an administrator shares with that

user. Administrators can additionally restrict individual connectors through custom permissions. Each connector supports three permissions: **Create and Update action**, **Share action**, and **Use action**. Custom permissions profiles apply at the user, role, or account level. The most specific level takes precedence. For more information, see [Custom permissions](#).

- **What to verify** – Before enabling a connector, verify the authentication method and OAuth scopes granted, whether the connector is shared broadly or narrowly, whether it is on-demand or automated, and what the service-level authorization permits. The credential you configure defines the connector's reach into the connected system. The connector can act on whatever the granted scopes and the service-level authorization permit, and nothing more.

For more information about action types, see [On-demand actions](#) and [Automated workflows](#). For details on authentication methods, see [Authentication methods](#).

Remote MCP servers

With remote Model Context Protocol (MCP) server integrations, the AI assistant can invoke tools that those servers provide. When the assistant invokes a tool, it sends tool invocation parameters derived from conversational context to the MCP server endpoint.

- **What triggers it** – The AI assistant selects a tool during a user-initiated or automated conversation. During connection, Amazon Quick discovers and registers the tools that the MCP server provides. After discovery completes, each tool is listed as an action that you can review and turn on.
- **Who controls it** – An administrator creates the MCP integration in the console. The integration specifies the endpoint URL, the connection type (public internet or private VPC), and the authentication method. You need an Enterprise subscription. Administrators can restrict individual MCP connectors through custom permissions. Each connector supports three permissions: **Create and Update action**, **Share action**, and **Use action**. For more information, see [Custom permissions](#).
- **Transport restriction** – MCP integration supports remote servers only. Local stdio connections are not supported. An MCP integration cannot invoke a process running on a local machine.
- **What to verify** – Before enabling an MCP integration, verify whether authentication is enabled (unauthenticated access is a selectable option), the OAuth scope if user authentication is selected, the network path (public internet or private VPC), and what tools the server exposes. Confirm that only the tools required for the intended use case are turned on.

For more information, see [Model Context Protocol \(MCP\) integration](#).

Web search

When web search is enabled, Amazon Quick sends queries derived from user requests to external web search providers. The assistant uses web search automatically in chat and research sessions to gather information from publicly available online sources.

- **What triggers it** – The assistant selects web search during a user-initiated chat or research session. You can turn off web search for individual queries.
- **Who controls it** – An administrator enables or disables web search for the organization. When web search is enabled, you can turn it off for each query. Web search queries route through designated AWS Regions based on the account's primary Region.
- **What to verify** – Before enabling web search, determine whether your use case requires external web data and understand how queries route across Regions.

For more information, see [Admin controls](#).

For more information about cross-Region routing for web search, see [Cross-Region calls for web search](#).

Third-party research providers

Amazon Quick Research integrates with third-party data providers such as FactSet, IDC, and S&P Global. When a user runs a research session, the agent determines which providers to use based on relevance to the research objective and sends queries to those providers.

- **What triggers it** – The research agent selects providers automatically based on relevance. Users choose which provider categories to include for each research session.
- **Who controls it** – Access to third-party research providers is layered across multiple controls. An administrator sets up each integration individually and shares it only with the users who need access. Unshared integrations are unavailable. Premium providers (FactSet, IDC, S&P Global Energy, S&P Global Market Intelligence) additionally require an administrator-provisioned license. If you don't provision a license, those providers stay unavailable. The Research capability in custom permissions removes access to Research – and therefore all third-party providers – for a given profile. Restricting the "Use internet to enhance results" capability separately removes web-based search in chat agents and Research.

- **What to verify** – Verify that each third-party data integration is shared only with authorized users. For premium providers, verify that licenses are provisioned only where needed. For public providers, verify that your acceptable-use policies permit sending research queries to external sources. Use the Research capability in custom permissions to remove Research entirely for profiles that must not reach any external provider. For more information about custom permissions, see [Custom permissions](#).

For more information, see [Using third party data in Amazon Quick Research](#) and [Select research materials](#).

Quick Apps and public access

Quick Apps run inside a security sandbox that blocks direct external network access. All communication with external systems goes through the secure bridge API or a registered connector. When an app owner enables public sharing, app content renders to anonymous internet viewers who have no identity, no integrations, and are rate-limited.

- **What triggers it** – The app owner creates and publishes an app. Anonymous viewers access it through a public URL.
- **Who controls it** – The app owner controls sharing and public access settings. Public access is available only to Free and Plus accounts. The sandbox is platform-enforced and blocks app code from making direct HTTP requests to external servers or loading external resources.
- **What to verify** – Before making an app public, verify that public viewers have no access to integrations, that usage counts against the owner's quota, and that the sandbox restrictions meet your security requirements.

For more information, see [Security and sandbox in apps in Quick](#).

Embedding

Embedding renders Amazon Quick dashboard visuals, analytics, and generative BI content in an external web application. Domain allowlisting controls where embedded content can render. Row-level security (RLS) and column-level security (CLS) are enforced for embedded viewers.

- **What triggers it** – An administrator or developer configures embedding through IAM API calls. End users then access the embedded content in the host application.

- **Who controls it** – An Enterprise subscription is required. The administrator or developer configures the AllowedDomains parameter to limit which domains can host embedded content. The IAM policy condition AllowedEmbeddingDomains provides additional restriction.
- **What to verify** – Before enabling embedding, verify that every domain in the allowlist is controlled by your organization, that RLS and CLS are configured for the underlying datasets, and that you have an Enterprise subscription.

For more information, see [Embedding Amazon Quick Sight analytics into your applications](#).

Exports and downloads

Users can export dashboard data as PDF or CSV files, download research reports, and export analysis data. Exported data goes to the user's local device or, for scheduled email reports, to the configured email delivery destination.

- **What triggers it** – A user initiates an export or download action in the console.
- **Who controls it** – An administrator can restrict export capabilities through custom permissions. For example, you can disable access to the export-to-PDF option on analyses and dashboards. An Enterprise subscription is required for custom permissions.
- **What to verify** – No single consolidated export-security configuration exists. Verify whether export capabilities are restricted through custom permissions for the relevant roles, whether scheduled email reports are limited to intended recipients, and whether your compliance requirements impose constraints on where exported data can be stored.

For more information about restricting capabilities, see [Custom permissions](#).

Log and metric delivery

Use log and metric delivery to send data from Amazon Quick to destinations that you control. Unlike the other paths on this page, data goes to your own AWS infrastructure rather than to a third party. Log types include chat conversations, user feedback, agent and research hours, index storage usage, and knowledge base file sync results.

Sensitive data in chat logs

Chat logs can contain sensitive or personally identifiable data.

- **What triggers it** – An administrator configures delivery. After configuration, delivery is continuous or scheduled.
- **Who controls it** – An administrator configures the delivery destination (CloudWatch Logs, Amazon S3, or Firehose). IAM permissions are required to set up delivery. When a destination uses a customer managed KMS key, the key policy must allow the delivery service principal.
- **What to verify** – Before configuring log delivery, verify the encryption and access policies at the destination, the retention policies at the destination, and that personnel with access to the destination understand that chat logs can contain sensitive data.

For more information, see [Monitoring Amazon Quick using CloudWatch Logs](#).

Cross-Region inference

Cross-Region inference is processing within a defined geography, not customer-enabled egress. It is included on this page to prevent confusion, because it involves data moving between Regions. However, it differs from every other path described on this page. There is no customer-chosen destination. No data reaches a third party. Data stays on the AWS network, and stored data remains in the primary Region.

- **What moves** – Inference inputs (prompts) and outputs (results) can move outside of the primary Region during processing. All data is transmitted encrypted across the AWS secure network.
- **What triggers it** – Amazon Quick automatically selects the optimal Region within a defined geography to process inference requests. You do not initiate or schedule this routing.
- **Who controls it** – Cross-Region inference is not a customer-configurable outbound path. Inference is processed within the Regions that make up your geography, no data reaches a third party, and stored data remains in your primary Region. Because there is no outbound destination, there is no destination for you to approve or restrict. You cannot disable cross-Region inference. Instead, verify which geography your Region belongs to. Confirm that the geography boundary satisfies your data residency requirements.
- **Boundaries** – Cross-Region inference requests stay within the AWS Regions that are part of the geography where the data originally resides. Stored data remains only in the primary Region.

For more information, see [Cross-Region inference for Australia, Japan, Europe, and the United States](#).

User lifecycle and data handling in Amazon Quick

What happens to a user's data and assets in Amazon Quick depends on how the user is removed. This page describes the three removal paths and their effect on a user's assets. It also covers the cleanup steps you must complete as an administrator.

Topics

- [How users are removed](#)
- [What is deleted immediately](#)
- [Assets with multiple owners](#)
- [Orphaned assets](#)
- [Deactivation compared with deletion](#)
- [Account termination](#)
- [Conversations and Amazon Q data](#)
- [When access ends](#)

How users are removed

How a user is removed determines what happens to the assets they owned. The following table describes the three removal paths.

Removal path	Access	What happens to the user's assets
Console deletion by an administrator	Ends when the user is deleted.	The administrator chooses whether to transfer the user's assets to another user or delete them. If no transfer user is chosen, assets owned solely by the deleted user are deleted, and assets that have other owners remain available to those owners with the deleted user removed from their permissions.

Removal path	Access	What happens to the user's assets
DeleteUser API operation	Ends when the user is deleted.	No asset cleanup runs. The user's assets are neither transferred nor deleted; they remain in the account without an owner until an administrator transfers them.
Removal from a directory group (IAM Identity Center or Active Directory)	The user can no longer sign in.	No asset cleanup runs automatically. Amazon Quick detects the removal and lists the user in the Inactive users list on the Manage users page, where an administrator can review them and apply the same transfer-or-delete choice as console deletion. Until an administrator acts, the user's assets remain in the account. This detection applies only to IAM Identity Center and Active Directory.

To complete cleanup, administrators use the asset management console to transfer assets owned by a departed user, and review users removed from IAM Identity Center or Active Directory on the **Manage users** page, where those users appear in the **Inactive users** list. For the console deletion procedure, see [Deleting a Amazon Quick user account](#). To transfer assets that no longer have an owner, see [Amazon Quick asset management](#). For users removed from IAM Identity Center or Active Directory, see [Deleting Enterprise accounts](#).

The Inactive users list applies to directory-federated users only

Amazon Quick detects removed users and populates the **Inactive users** list only for accounts that federate through IAM Identity Center or Active Directory. Users that you create directly in Amazon Quick, including IAM federation users, are managed at the user

level, so Amazon Quick does not detect their removal from an external identity provider. Remove these users with console deletion or the `DeleteUser` API operation, and transfer or delete their assets as part of that removal.

If your knowledge bases use document-level access control lists (ACLs), also update the ACL configuration to remove the departing user, and refresh the affected knowledge bases. User removal does not update knowledge base ACLs. For more information, see [Best practices for managing ACLs in knowledge bases](#).

What is deleted immediately

When you delete a user through any path, Amazon Quick permanently deletes the user's metadata, including their profile and role membership.

Assets with multiple owners

Amazon Quick never deletes an asset that has more than one owner when a user is removed. Amazon Quick removes the departed user from the asset's permissions, and the asset remains available to its other owners.

Orphaned assets

No user can access assets that remain in the account without an owner – for example, after deletion with the `DeleteUser` API operation, or after identity-provider removal that is pending administrator review – until you transfer them. Use the asset management console to find and transfer assets that a departed user owned. For more information, see [Amazon Quick asset management](#).

Deactivation compared with deletion

Deactivating a user revokes the user's access but leaves the user's assets associated with them. Deactivation occurs, for example, when you remove the user from the IAM Identity Center or Active Directory group that grants Amazon Quick access. Reactivating the user restores their access to those assets. Deletion is permanent.

A user who is removed from IAM Identity Center or Active Directory appears in the **Inactive users** list on the **Manage users** page and remains there until an administrator removes them.

A user who is deleted but not deactivated can sign in again as a new user, without access to their previous assets.

For more information, see [Deactivating user accounts](#).

Account termination

Unsubscribing from Amazon Quick permanently deletes all of your Amazon Quick data, including all metadata, any data in SPICE, and Amazon Q data. For more information about account settings, see [Account details in Amazon Quick](#). For how unsubscribing affects the Amazon Q data key, see [Amazon Q data key](#).

Conversations and Amazon Q data

A user's chat conversations are private to that user. When you delete or deprovision a user, Amazon Quick does not transfer their conversation history to any other user, and you can't access it through the application.

Deleting a user does not immediately delete the user's conversation history. Amazon Quick retains conversations, messages, and chat file attachments on a rolling 30-day window measured from activity, and then permanently deletes them. Users can delete their own conversations at any time. When you unsubscribe from Amazon Quick, Amazon Quick permanently deletes all Amazon Q data, including conversation history.

Amazon Quick does not delete feedback that a user submitted when that user is deleted. Amazon Quick retains analytics data according to your instance's data retention policies.

When access ends

Removing or deleting a user prevents new sign-ins to Amazon Quick. Amazon Quick also denies requests from the user's existing sessions after the change takes effect, which can take up to five minutes.

Incident response, logging, and monitoring in Amazon Quick

Intended audience: System administrators and Amazon Quick administrators

Amazon Quick provides multiple monitoring and audit signals that address different aspects of security, operations, and compliance. Each signal has distinct coverage, latency, retention, and access-control characteristics.

- **CloudTrail** – Records supported Amazon Quick and Amazon Quick Sight API operations and a documented set of non-API events, such as dashboard views and user-management actions. For chat conversations and feedback, use CloudWatch vended logs.
- **CloudWatch vended logs** – Deliver chat conversations, user feedback, agent hours usage, index storage usage, and knowledge base file sync results to destinations that you control. Configure vended log delivery shortly after enabling Amazon Quick.
- **CloudWatch metrics** – Provide near-real-time operational metrics and support CloudWatch alarms.
- **Amazon Quick analytics** – Provide usage, adoption, feedback, and selected security-related insights to IAM administrators.
- **Feature-specific reports** – Provide operational detail for the feature that produces them, such as knowledge base sync reports.
- **Amazon EventBridge** – CloudTrail events route to Amazon EventBridge on a best-effort basis. You can create rules that match Amazon Quick events and route them to targets such as Lambda functions, Amazon SNS topics, or Amazon SQS queues for automated response.

Use these sources together when you design monitoring and incident-response procedures. Before you rely on a signal for compliance, detection, or investigation, confirm that it covers the event you intend to track.

The following table helps you choose the right signal for your monitoring need.

To answer	Use	More information
Who performed an administrative or API action, from where, and when	CloudTrail	Monitoring Amazon Quick using CloudTrail
Who viewed a dashboard, or which non-API events occurred	CloudTrail non-API events	Tracking non-API events by using CloudTrail logs
Which API calls occurred for AI features (flows, agents, automations, connectors)	CloudTrail (management and data events)	Monitoring Amazon Quick using CloudTrail; Logging Amazon Quick data events in CloudTrail

To answer	Use	More information
What users asked and what Quick answered	CloudWatch vended logs (CHAT_LOGS)	Monitoring Amazon Quick using CloudWatch Logs
How users rated responses and why	Vended logs (FEEDBACK_LOGS) or analytics	Monitoring Amazon Quick using CloudWatch Logs ; Using the Amazon Quick analytics dashboard
Whether a document synced into a knowledge base, and why it failed or was skipped	Vended logs (KB_FILE_SYNC_LOGS) or console sync reports	Monitoring Amazon Quick using CloudWatch Logs ; Sync reports and observability
Whether a specific user can access a specific synced document	Sync report ACL verification	Sync reports and observability
Index storage per knowledge base or Space	Vended logs (INDEX_USAGE_LOGS) or CloudWatch metrics	Monitoring Amazon Quick using CloudWatch Logs ; Monitoring Amazon Quick using CloudWatch metrics
Operational health: load times, ingestion failures, connector errors, SPICE capacity	CloudWatch metrics and alarms	Monitoring Amazon Quick using CloudWatch metrics
Adoption, engagement, feedback trends, agent-hours consumption	Analytics dashboard	Using the Amazon Quick analytics dashboard

Use the following checklist to configure monitoring for your environment:

1. Create a CloudTrail trail for the AWS accounts and AWS Regions that require audit logging.
2. Configure CloudWatch vended log delivery shortly after enabling Amazon Quick AI features.

3. When vended-log destinations use a customer managed AWS KMS key, allow `delivery.logs.amazonaws.com` in the key policy.
4. Chat logs can contain sensitive or personally identifiable data. Filter this information at subscription setup, or apply CloudWatch Logs data-protection masking policies.
5. Configure CloudWatch alarms on operational metrics and map them to your incident-response procedures.
6. Grant `quicksight:QuickSuiteUsageMetrics` (analytics access) only to authorized administrators.
7. Revalidate your monitoring design when you enable a new Amazon Quick capability.

Topics

- [Monitoring Amazon Quick using CloudWatch Logs](#)
- [Monitoring Amazon Quick using CloudTrail](#)
- [Monitoring Amazon Quick using CloudWatch metrics](#)

Monitoring Amazon Quick using CloudWatch Logs

You can use [Amazon CloudWatch Logs](#) to deliver chat conversations, user feedback, agent hours usage, and index storage usage in Amazon Quick for you to analyze. These logs can be delivered to multiple destinations, such as CloudWatch, Amazon S3, or Amazon Data Firehose (standard rates apply).

Important

Set up vended log delivery shortly after enabling Amazon Quick AI features. Logs are not retroactive – you only receive events that occur after delivery is configured.

The following are examples of tasks you can complete with logs from Amazon Quick:

- Identify common user queries and pain points by reviewing the chat message content.
- Monitor the quality of responses by looking at metrics like `feedbackReason`.
- Understand user sentiment and satisfaction by analyzing the feedback data, including comments and usefulness ratings.
- Generate custom dashboards and reports to track key metrics and trends over time.

- Identify and analyze cases where the chat returned no answer or the user query was blocked.
- Monitor agent hours usage.
- Track index storage usage across knowledge bases and Spaces.
- Track knowledge base sync operations and identify document crawl or indexing failures.

Important

Logs from conversations might include sensitive or personally identifiable data passed in the chats. You can filter out this information from your logs when setting up logs subscription. Or you can mask this data on your logs using CloudWatch Logs masking policies. For more information, see [Help protect sensitive log data with masking](#).

Supported log destinations

Amazon Quick can deliver logs to the following destinations:

- **Amazon CloudWatch Logs** – For real-time monitoring and analysis
- **Amazon S3** – For long-term storage and batch processing
- **Amazon Data Firehose** – For streaming analytics and data transformation

Supported log types

Amazon Quick supports the following log types. You specify one of these values as the `logType` when you create a delivery source.

- **CHAT_LOGS** – Chat conversation interactions, including the user message, system response, and status.
- **FEEDBACK_LOGS** – User feedback on chat responses, such as usefulness ratings and reasons.
- **AGENT_HOURS_LOGS** – Agent hours usage for each agent.
- **AGENT_METADATA_LOGS** – Lifecycle events for chat agents, such as creation, updates, and deletion.
- **INDEX_USAGE_LOGS** – Per-source index storage metrics for knowledge bases and Spaces.
- **KB_FILE_SYNC_LOGS** – Per-document knowledge base sync status, including crawl and indexing failures.

- DLP_LOGS – Data loss prevention enforcement decisions and configuration changes.

Prerequisites

Before you can enable logging, make sure that you have the following:

- An active Amazon Quick instance with Enterprise or Professional subscriptions
- Appropriate IAM permissions to configure log delivery
- A destination configured for your logs (CloudWatch Logs, Amazon S3 bucket, or Firehose)

Configure logging

To enable logging for Amazon Quick chat and feedback, configure IAM permissions, create a delivery source and destination, and verify log delivery.

Topics

- [Set up IAM permissions](#)
- [Configure log subscription](#)
- [Verify log delivery](#)

Set up IAM permissions

To set up CloudWatch Logs for Amazon Quick, use the following IAM policy examples to grant the necessary permissions.

```
{
  "Version": "2012-10-17"
  ,
  "Statement": [{
    "Sid": "QuicksightLogDeliveryPermissions",
    "Effect": "Allow",
    "Action": "quicksight:AllowVendedLogDeliveryForResource",
    "Resource": "arn:aws:quicksight:region:account-id:account/account-id"
  }]
}
```

You must also allow the `delivery.logs.amazonaws.com` service principal in your customer managed AWS KMS key policy.

```
{
  "Effect": "Allow",
  "Principal": {
    "Service": "delivery.logs.amazonaws.com"
  },
  "Action": [
    "kms:GenerateDataKey",
    "kms:Decrypt"
  ],
  "Resource": "*",
  "Condition": {
    "StringLike": {
      "kms:EncryptionContext:SourceArn": "arn:partition:logs:region:account-id:*"
    }
  }
}
```

Configure log subscription

For example IAM policies with all the required permissions for your specific logging destination, see [Enable logging from AWS services](#) in the *Amazon CloudWatch Logs User Guide*.

Create a delivery source with the [PutDeliverySource](#) CloudWatch Logs API operation. Give the delivery source a name and for `resourceArn`, specify the ARN of your application. For `logType`, specify one of the supported log types (see [the section called “Supported log types”](#)).

```
{
  "logType": "CHAT_LOGS",
  "name": "my-quick-suite-delivery-source",
  "resourceArn": "arn:aws:quicksight:your-region:your-account-id:account/account-id"
}
```

```
{
  "logType": "FEEDBACK_LOGS",
  "name": "my-quick-suite-delivery-source",
  "resourceArn": "arn:aws:quicksight:your-region:your-account-id:account/account-id"
}
```

```
{
  "logType": "AGENT_HOURS_LOGS",
  "name": "my-quick-suite-delivery-source",
```



```
"resourceArn": "arn:aws:quicksight:your-region:your-account-id:account/account-id"
}
```

```
{
  "logType": "INDEX_USAGE_LOGS",
  "name": "my-quick-index-usage-delivery-source",
  "resourceArn": "arn:aws:quicksight:your-region:your-account-id:account/account-id"
}
```

```
{
  "logType": "KB_FILE_SYNC_LOGS",
  "name": "my-quick-kb-file-sync-delivery-source",
  "resourceArn": "arn:aws:quicksight:your-region:your-account-id:account/account-id"
}
```

```
{
  "logType": "AGENT_METADATA_LOGS",
  "name": "my-quick-agent-metadata-delivery-source",
  "resourceArn": "arn:aws:quicksight:your-region:your-account-id:account/account-id"
}
```

```
{
  "logType": "DLP_LOGS",
  "name": "my-quick-dlp-delivery-source",
  "resourceArn": "arn:aws:quicksight:your-region:your-account-id:account/account-id"
}
```

To enable user conversation logging with the CloudWatch Logs API operations, you call the `PutDeliverySource`, `PutDeliveryDestination`, and `CreateDelivery` API operations.

Note

Vended-log delivery is per Region. Each delivery captures only the activity that occurs in the Region specified by the `resourceArn` in `PutDeliverySource`. To receive logs from multiple Regions, configure a separate delivery in each Region.

Verify log delivery

After you configure log delivery, verify that logs are being delivered to your destination:

- **Verify the setup:** Verify the list of deliveries that have been created in the account by using the `DescribeDeliveries` API in CloudWatch Logs.
- **CloudWatch Logs:** Check the specified log group for new log streams.
- **Amazon S3:** Monitor your bucket for new log files.
- **Firehose:** Verify data is flowing through your delivery stream.

Log schema and format

Each Amazon Quick log type has its own schema. The field names, including the log type and account identifiers, vary by log type – for example, chat and feedback logs use `logType` and `accountId`, while other log types use `log_type` and `account_id`. The fields for each log type are listed in the following sections.

Chat logs

Chat logs capture conversation interactions:

Note

Temporary conversations that are excluded from history and memory are still delivered to chat logs. Each record includes the `user_arn` and `user_type` of the user who initiated the conversation.

- `user_arn` – Amazon Quick user ARN associated with the event
- `user_type` – Amazon Quick user type associated with the event (for example, `ADMIN_PRO`)
- `status_code` – Status of the chat request (for example, `success`, `request_blocked`, `no_answer_found`)
- `conversation_id` – Unique ID for the user conversation
- `system_message_id` – System-generated message ID
- `message_scope` – Scope of the message (for example, `all_resources`, `specific_resources`, `no_resources`)
- `user_message_id` – Unique ID of the user message
- `user_message` – User message in the conversation

- `agent_id` – Unique ID of the chat agent, or `SYSTEM` for the default agent
- `flow_id` – Unique ID of the Amazon Quick Flow, or `-` if not a flow invocation
- `system_text_message` – System response in the conversation
- `user_selected_resources` – List of resources selected by user
- `action_connectors` – List of connectors available in the conversation
- `cited_resource` – List of cited resources
- `file_attachment` – List of files attached by user
- `resource_arn` – Resource ARN of your Amazon Quick account
- `event_timestamp` – Timestamp of the event
- `logType` – `CHAT_LOGS`
- `accountId` – AWS account ID

The following is an example of chat logs:

```
{
  "user_arn": "arn:aws:quicksight:us-east-1:111122223333:user/default/johndoe",
  "user_type": "ADMIN_PRO",
  "status_code": "success",
  "conversation_id": "c11ba72c-ff18-4213-9686-1952bb547c19",
  "system_message_id": "42a37690-1804-442b-8368-3d34570dd2cd",
  "message_scope": "all_resources",
  "user_message_id": "5ec45e03-bf22-40c7-b32a-d69eb015f86b",
  "user_message": "What is in this document?",
  "agent_id": "SYSTEM",
  "flow_id": "-",
  "system_text_message": "Here is a summary of the attached document.",
  "user_selected_resources": [{"resourceId": "ALL", "resourceType": "space"}],
  "action_connectors": [{"actionConnectorId": "quicksuite-documentation"},
{"actionConnectorId": "quicksuite-websearch"}],
  "cited_resource": [{"citedResourceType": "document", "citedResourceId": null,
"citedResourceName": "example-document.html"}],
  "file_attachment": [{"fileAttachmentType": "html", "fileAttachmentName": "example-
document.html"}],
  "resource_arn": "arn:aws:quicksight:us-east-1:111122223333:account/111122223333",
  "event_timestamp": 1787185484508,
  "logType": "CHAT_LOGS",
  "accountId": "111122223333"
```

```
}
```

Feedback logs

Feedback logs capture user feedback on chat responses:

- `user_arn` – Amazon Quick user ARN associated with the event
- `user_type` – Amazon Quick user type associated with the event (for example, `ADMIN_PRO`)
- `status_code` – Status of the event delivery
- `conversation_id` – Unique ID of the conversation
- `system_message_id` – System generated message ID
- `user_message_id` – Unique ID of user message
- `feedback_type` – Type of feedback (for example, `Not Useful`, `Useful`)
- `feedback_reason` – Feedback reason selected by the user
- `feedback_details` – (Optional) Additional details provided by the user
- `resource_arn` – Resource ARN of your Amazon Quick account
- `event_timestamp` – Timestamp of the event
- `logType` – `FEEDBACK_LOGS`
- `accountId` – AWS account ID

The following is an example of feedback logs:

```
{
  "user_arn": "arn:aws:quicksight:us-east-1:111122223333:user/default/johndoe",
  "user_type": "ADMIN_PRO",
  "status_code": "success",
  "conversation_id": "c11ba72c-ff18-4213-9686-1952bb547c19",
  "system_message_id": "42a37690-1804-442b-8368-3d34570dd2cd",
  "user_message_id": "5ec45e03-bf22-40c7-b32a-d69eb015f86b",
  "feedback_type": "Not Useful",
  "feedback_reason": "Too wordy",
  "feedback_details": "The answer included too much background information.",
  "resource_arn": "arn:aws:quicksight:us-east-1:111122223333:account/111122223333",
  "event_timestamp": 1787185500123,
  "logType": "FEEDBACK_LOGS",
  "accountId": "111122223333"
```

```
}
```

Agent hours logs

This log type captures the usage logs for different agents within your Quick account:

- `user_arn` – Amazon Quick user ARN associated with the event
- `subscription_type` – Subscription tier of the user. Values: ENTERPRISE, PROFESSIONAL.
- `reporting_service` – The Quick surface that consumed agent hours. Current values include FLOW, AUTOMATION, and RESEARCH. New values might appear as additional Quick features begin metering agent hours.
- `usage_group` – Whether the usage is covered by the subscription entitlement or billed as overage. Values:
 - Included – Usage within the daily entitlement grant for the subscription tier. No incremental charge.
 - Extra – Overage beyond the daily grant. Billed on consumption.
- `usage_hours` – Decimal value indicating the agent hours consumed for this record
- `service_resource_arn` – ARN of the resource that consumed the hours (for example, a flow, automation, or research session)
- `resource_arn` – Resource ARN of your Amazon Quick account
- `event_timestamp` – Timestamp of the event
- `logType` – AGENT_HOURS_LOGS
- `accountId` – AWS account ID

The following is an example of Agent Hours logs:

```
{
  "user_arn": "arn:aws:quicksight:us-east-1:111122223333:user/default/johndoe",
  "subscription_type": "ENTERPRISE",
  "reporting_service": "RESEARCH",
  "usage_group": "Included",
  "usage_hours": 0.0928,
  "service_resource_arn": "arn:aws:quicksight::111122223333:research/a11b2bbc-
c123-3abc-a12b-12a34b5c678d",
  "resource_arn": "arn:aws:quicksight:us-east-1:111122223333:account/111122223333",
  "event_timestamp": 1787183715000,
  "logType": "AGENT_HOURS_LOGS",
}
```

```
"accountId": "111122223333"  
}
```

Agent metadata logs

Agent metadata logs capture lifecycle events for chat agents, including creation, updates, permission changes, and deletion. One record is emitted per agent operation:

- `user_arn` – Amazon Quick user ARN associated with the event.
- `event_name` – The agent lifecycle operation. Values include `CreateAgent`, `UpdateAgent`, `DeleteAgent`, `UpdateAgentPermissions`, and others as new operations are added.
- `event_version` – Schema version of the log record.
- `agent_id` – UUID of the agent.
- `agent_arn` – Full ARN of the agent.
- `agent_name` – Display name of the agent.
- `agent_status` – Current status of the agent (for example, `ACTIVE`).
- `request_id` – Request identifier for the operation.
- `description` – Agent description text.
- `spaces` – JSON array of spaces attached to the agent.
- `permissions_granted` – Permissions added in this operation.
- `permissions_revoked` – Permissions removed in this operation.
- `permissions_state` – Current permissions after the operation.
- `update_action` – The update action performed.
- `version` – Agent version number.
- `icon_id` – Icon identifier for the agent.
- `magic_builder_query` – The natural language query used to create the agent via the builder.
- `instructions` – Agent instructions.
- `failed_to_add_spaces` – Spaces that could not be attached.
- `failed_to_remove_spaces` – Spaces that could not be detached.
- `draft_discarded` – Whether a draft was discarded.
- `custom_prompt_input` – Custom prompt input configured for the agent.
- `welcome_message` – Welcome message shown to users of the agent.
- `starter_prompts` – Starter prompts suggested to users of the agent.

- `resource_arn` – Resource ARN of your Amazon Quick account.
- `event_timestamp` – Timestamp of the event.
- `log_type` – `AGENT_METADATA_LOGS`.
- `account_id` – AWS account ID.

Note

The `custom_prompt_input`, `welcome_message`, and `starter_prompts` fields contain sensitive content. When you configure a customer managed AWS KMS key for delivery, these fields are encrypted. Without a key, they are delivered in cleartext.

The following is an example of an agent metadata log for a `CreateAgent` event:

```
{
  "user_arn": "arn:aws:quicksight:us-east-1:111122223333:user/default/johndoe",
  "event_version": "1.0",
  "event_name": "CreateAgent",
  "agent_id": "12345678-90ab-cdef-1234-567890abcdef",
  "request_id": "fedcba98-7654-3210-fedc-ba9876543210",
  "agent_arn": "arn:aws:quicksight:us-east-1:111122223333:agent/12345678-90ab-cdef-1234-567890abcdef",
  "agent_name": "Project Status Assistant",
  "agent_status": "ACTIVE",
  "description": "An agent that helps track project status and find relevant documentation.",
  "spaces": "[{\"Arn\": \"arn:aws:quicksight:us-east-1:111122223333:space/11112222-3333-4444-5555-666677778888\"}]",
  "resource_arn": "arn:aws:quicksight:us-east-1:111122223333:account/111122223333",
  "event_timestamp": 1787185396000,
  "log_type": "AGENT_METADATA_LOGS",
  "account_id": "111122223333"
}
```

Index usage logs

Index usage logs capture per-source storage metrics for knowledge bases and Spaces. Events are published whenever a change occurs (created, updated, synced, or deleted):

- `user_arn` – Amazon Quick user ARN associated with the event

- `consumed_index_size` – Total size (in bytes) consumed by the entire index
- `source_type` – Type of source: SPACE or KB
- `source_name` – Display name of the Space or knowledge base
- `source_arn` – Full ARN of the source
- `consumed_source_size` – Size (in bytes) consumed by this individual source
- `consumed_source_doc_count` – Number of documents in this source
- `resource_arn` – Resource ARN of your Amazon Quick account
- `event_timestamp` – Timestamp of the event
- `log_type` – INDEX_USAGE_LOGS
- `account_id` – AWS account ID

The following is an example of index usage logs:

```
{
  "account_id": "111122223333",
  "event_timestamp": 1774911984257,
  "log_type": "INDEX_USAGE_LOGS",
  "user_arn": "arn:aws:quicksight::111122223333:user/default/user",
  "resource_arn": "arn:aws:quicksight:us-west-2:111122223333:account/111122223333",
  "consumed_index_size": 500000,
  "source_type": "SPACE",
  "source_name": "my-space",
  "source_arn": "arn:aws:quicksight:us-west-2:111122223333:space/2744af89-31b2-423b-93a2-69b0cd0d7fa1",
  "consumed_source_size": 244436,
  "consumed_source_doc_count": 2
}
```

 Note

Events are published per source on change. Not all sources emit events every day. To reconstruct the current state, use the most recent event per `source_arn`.

For information about building dashboards and running queries against index usage logs, see [the section called “Monitor index storage usage”](#).

Knowledge base file sync logs

Knowledge base file sync logs capture per-document sync status events. One log record is emitted per document per sync run:

- `document_id` – Original document identifier such as a URL or file path.
- `document_title` – Document title.
- `document_status` – Terminal document status. Values: ADDED, MODIFIED, UNMODIFIED, DELETED, SKIPPED, FAILED.
- `sync_result` – High-level availability result. Values: AVAILABLE or UNAVAILABLE.
- `sync_id` – Sync job execution ID.
- `data_source_id` – Identifier of the data source that the knowledge base is connected to.
- `source_uri` – Source URL of the document.
- `error_message` – Error description when status is FAILED or SKIPPED.
- `error_mitigation` – Actionable guidance for resolving the error.
- `error_type` – Error code when status is FAILED or SKIPPED.
- `knowledge_base_id` – UUID of the knowledge base that produced the log.

The following table describes the `document_status` values and their corresponding `sync_result`.

<code>document_status</code>	<code>sync_result</code>	Meaning
ADDED	AVAILABLE	New document successfully indexed
MODIFIED	AVAILABLE	Existing document re-indexed with changes
UNMODIFIED	AVAILABLE	Document content unchanged, no re-indexing needed
DELETED	UNAVAILABLE	Document removed from index
SKIPPED	UNAVAILABLE	Document filtered during crawl, such as by robots.txt or size limit
FAILED	UNAVAILABLE	Document failed during crawl or indexing

The following example shows a successful knowledge base file sync log where a document was added to the knowledge base:

```
{
  "resource_arn": "arn:aws:quicksight:us-west-2:111122223333:account/111122223333",
  "event_timestamp": 1781296858575,
  "log_type": "KB_FILE_SYNC_LOGS",
  "account_id": "111122223333",
  "document_id": "https://docs.aws.amazon.com/quick/latest/userguide/monitoring-
cloudwatch-logs.html",
  "document_title": "Monitoring Amazon QuickSight usage using CloudWatch Logs",
  "document_status": "ADDED",
  "sync_result": "AVAILABLE",
  "sync_id": "86a70a9a-cad9-4fc6-8881-e3909c8954d2",
  "data_source_id": "56225744-18bc-4373-a91f-861dd1c3d566",
  "source_uri": "https://docs.aws.amazon.com/quick/latest/userguide/monitoring-
cloudwatch-logs.html",
  "knowledge_base_id": "b0bd0a47-8095-439d-9dff-c64bd5fe3fa3"
}
```

The following example shows a skipped document log where crawling was skipped due to a validation error:

```
{
  "resource_arn": "arn:aws:quicksight:us-west-2:111122223333:account/111122223333",
  "event_timestamp": 1781296492951,
  "log_type": "KB_FILE_SYNC_LOGS",
  "account_id": "111122223333",
  "document_id": "https://docs.aws.amazon.com/quick/latest/userguide/blocked-
page.html",
  "document_title": "Blocked Page",
  "document_status": "SKIPPED",
  "sync_result": "UNAVAILABLE",
  "sync_id": "86a70a9a-cad9-4fc6-8881-e3909c8954d2",
  "data_source_id": "56225744-18bc-4373-a91f-861dd1c3d566",
  "source_uri": "https://docs.aws.amazon.com/quick/latest/userguide/blocked-
page.html",
  "error_message": "This URL wasn't crawled because crawling isn't allowed by its
robots.txt file.",
  "error_mitigation": "Contact the website administrator for assistance.",
  "error_type": "VALIDATION_ERROR",
  "knowledge_base_id": "b0bd0a47-8095-439d-9dff-c64bd5fe3fa3"
}
```

```
}
```

DLP logs

Data loss prevention (DLP) logs capture DLP activity: the enforcement decision for each scanned file, and changes to your DLP configurations. Use them to audit policy changes and to monitor enforcement. The specific event is identified by `event_type`.

Every `DLP_LOGS` record includes these common fields:

- `resource_arn` – Resource ARN of your Amazon Quick account
- `event_timestamp` – Unix epoch time in milliseconds
- `log_type` – `DLP_LOGS`
- `account_id` – AWS account ID
- `event_type` – The DLP event: `DLP_FILE_BLOCKED`, `DLP_FILE_WARNED`, `DLP_INSPECTION_FAILED`, `DLP_SETTING_CREATED`, `DLP_SETTING_UPDATED`, or `DLP_SETTING_DELETED`
- `request_id` – Unique identifier for the event
- `user_arn` – Amazon Quick user ARN associated with the event (system for service-initiated events)

Enforcement events (`DLP_FILE_BLOCKED`, `DLP_FILE_WARNED`, `DLP_INSPECTION_FAILED`) add these fields:

- `dlp_job_id` – Unique ID of the DLP scan job
- `dlp_setting_id` – ID of the DLP configuration that evaluated the file
- `policy_action` – Enforcement action applied (`BLOCK`, `WARN`)
- `file_name` – Name of the scanned file (customer content; see the note that follows)
- `file_size` – Size of the scanned file, in bytes
- `failure_type` – (`DLP_INSPECTION_FAILED` only) Category of the failure, for example `CUSTOMER_ERROR`
- `policy_message` – (`DLP_FILE_WARNED` and `DLP_INSPECTION_FAILED`) The warning message shown to the user, or the reason the inspection failed

Configuration events (DLP_SETTING_CREATED, DLP_SETTING_UPDATED, DLP_SETTING_DELETED) add these fields:

- `dlp_setting_id` – ID of the DLP configuration
- `dlp_setting_name` – Display name of the configuration
- `status` – Configuration status (ACTIVE, INACTIVE) – reflects whether enforcement is enabled or disabled
- `provider_type` – DLP provider (for example, MICROSOFT_PURVIEW)
- `auth_type` – Provider authentication type (for example, CLIENT_SECRET)
- `provider_outage_mode` – Action applied when the provider is unavailable (for example, BLOCK)
- `unmapped_action` – Default action for files with no mapped label (for example, BLOCK)
- `last_updated_by` – ARN of the principal that made the change
- `changes` – (DLP_SETTING_UPDATED only) The fields that changed

The following is an example of a blocked file event:

```
{
  "resource_arn": "arn:aws:quicksight:us-east-1:111122223333:account/111122223333",
  "event_timestamp": 1786510225000,
  "log_type": "DLP_LOGS",
  "account_id": "111122223333",
  "event_type": "DLP_FILE_BLOCKED",
  "request_id": "422518600001076813548728842",
  "dlp_setting_id": "dlp-config-prod-01",
  "dlp_job_id": "dlpjob-internal-1c5e20d9-8af4-4ad9-95ee-18c34cd62c76",
  "policy_action": "BLOCK",
  "file_name": "confidential.docx",
  "file_size": 64,
  "user_arn": "arn:aws:quicksight:us-east-1:111122223333:user/default/johndoe"
}
```

The following is an example of an inspection failed event:

```
{
  "resource_arn": "arn:aws:quicksight:us-east-1:111122223333:account/111122223333",
  "event_timestamp": 1786488697000,
  "log_type": "DLP_LOGS",
```

```

"account_id": "111122223333",
"event_type": "DLP_INSPECTION_FAILED",
"request_id": "421276000001518855116675749",
"dlp_setting_id": "dlp-config-prod-01",
"dlp_job_id": "dlpjob-internal-f25e8e10-d831-4479-8f0d-e32d7b12b123",
"policy_action": "BLOCK",
"failure_type": "CUSTOMER_ERROR",
"policy_message": "The QuickSight service role for SECRETS_MANAGER access is not
configured. Verify your account's QuickSight service role setup.",
"file_name": "confidential.docx",
"file_size": 64,
"user_arn": "system"
}

```

The following is an example of a configuration updated event:

```

{
  "resource_arn": "arn:aws:quicksight:us-east-1:111122223333:account/111122223333",
  "event_timestamp": 1786492367000,
  "log_type": "DLP_LOGS",
  "account_id": "111122223333",
  "event_type": "DLP_SETTING_UPDATED",
  "request_id": "327788200001887837483925731",
  "dlp_setting_id": "dlp-config-prod-01",
  "dlp_setting_name": "Production DLP policy",
  "status": "ACTIVE",
  "provider_type": "MICROSOFT_PURVIEW",
  "auth_type": "CLIENT_SECRET",
  "provider_outage_mode": "BLOCK",
  "unmapped_action": "BLOCK",
  "last_updated_by": "arn:aws:sts::111122223333:assumed-role/DlpAdmin/jane",
  "changes": { "unmapped_action": { "old": "WARN", "new": "BLOCK" } },
  "user_arn": "system"
}

```

Note

The `file_name` field is customer content. When you configure a customer managed AWS KMS key for delivery, this field is encrypted. Without a key, Quick delivers it in cleartext.

Security considerations

- **Encryption** – Use customer-managed AWS KMS keys for sensitive data.
- **Access control** – Implement least-privilege IAM policies.
- **Data retention** – Configure appropriate retention policies for your compliance requirements.

Monitoring Amazon Quick using CloudTrail

Amazon Quick is integrated with AWS CloudTrail. This service provides a record of actions taken by a user, role, or an AWS service in Amazon Quick. CloudTrail captures API calls for Amazon Quick as events. The calls captured include some calls from the console and all code calls to API operations. If you create a trail, you can enable continuous delivery of CloudTrail events to an Amazon S3 bucket, including events for Amazon Quick. If you don't configure a trail, you can still view the most recent events in the CloudTrail console in **Event history**. Using the information collected by CloudTrail, you can determine who made a request to Amazon Quick and when. You can also identify the source IP address, the specific request, and other details.

Note

CloudTrail records supported Amazon Quick API operations and the non-API events documented on this page. When you examine event samples, event records include selected metadata such as account ID, resource ID, and user identity. For chat conversations and feedback, use [Monitoring Amazon Quick using CloudWatch Logs](#). For the full signal-by-capability model, see [Incident response, logging, and monitoring in Amazon Quick](#).

Amazon Quick publishes operational metrics to CloudWatch, and you can create CloudWatch alarms on those metrics. For more information, see [Creating alarms with the Amazon CloudWatch console](#). To alert on individual CloudTrail events, you can develop a custom solution that processes CloudTrail logs.

Amazon Quick service status can be viewed on the [Service Health Dashboard](#).

By default, CloudTrail encrypts the log files it delivers to your bucket using [Amazon S3-managed encryption keys \(SSE-S3\)](#). For a directly manageable security layer, you can instead use [AWS KMS-managed keys \(SSE-KMS\)](#) for your CloudTrail log files. SSE-KMS encrypts the log files but not the digest files. Digest files are encrypted with [Amazon S3-managed encryption keys \(SSE-S3\)](#).

To learn more about CloudTrail, including how to configure and enable it, see the [AWS CloudTrail User Guide](#).

Topics

- [Logging Amazon Quick operations with AWS CloudTrail](#)
- [Discovering Amazon Quick events in your trail](#)
- [Example: Amazon Quick log file entries](#)
- [Logging Amazon Quick data events in CloudTrail](#)
- [Tracking non-API events by using CloudTrail logs](#)

Logging Amazon Quick operations with AWS CloudTrail

Intended audience: System administrators

CloudTrail is enabled on your AWS account when you create the account. When supported event activity occurs in Amazon Quick, that activity is recorded in a CloudTrail event along with other AWS service events in **Event history**. You can view, search, and download recent events in your AWS account. For more information, see [Viewing Events with CloudTrail Event History](#).

For an ongoing record of events in your AWS account, including events for Amazon Quick, create a trail. A *trail* enables CloudTrail to deliver log files to an Amazon S3 bucket. By default, when you create a trail in the console, the trail applies to all . The trail logs events from all Regions in the AWS partition and delivers the log files to the Amazon S3 bucket that you specify. Additionally, you can configure other AWS services to further analyze and act upon the event data collected in CloudTrail logs. For more information, see the following:

- [Overview for Creating a Trail](#)
- [CloudTrail Supported Services and Integrations](#)
- [Configuring Amazon SNS Notifications for CloudTrail](#)
- [Receiving CloudTrail Log Files from Multiple Regions](#) and [Receiving CloudTrail Log Files from Multiple Accounts](#)
- [Cross-Account CloudTrail Logging](#) in the AWS Lake Formation Developer Guide – This topic includes instructions for including principal identities in cross-account CloudTrail logs.

Amazon Quick supports logging the following actions as events in CloudTrail log files:

- Whether the request was made with root or AWS Identity and Access Management user credentials
- Whether the request was made with temporary security credentials for an IAM role or federated user
- Whether the request was made by another AWS service

For more information on user identity, see the [CloudTrail userIdentity Element](#).

By default, each Amazon Quick log entry contains the following information:

- **userIdentity** – User identity
- **eventTime** – Event time
- **eventID** – Event ID
- **readOnly** – Read only
- **awsRegion** – AWS Region
- **eventSource (quicksight)** – Source of the event (Amazon Quick)
- **eventType (AwsServiceEvent)** – Event type (AWS service event)
- **recipientAccountId (customer AWS account)** – Recipient account ID (Customer AWS account)

Some events include a `userIdentity.invokedBy` field set to `quicksight.amazonaws.com`. This indicates that the event was triggered by the Amazon Quick service on behalf of the account rather than by a direct user action. You can use this field to distinguish user-initiated operations from service-internal calls when analyzing your trail.

 Note

CloudTrail displays the user identity as unknown for users that were provisioned directly by Amazon Quick (legacy provisioning). Users authenticated through IAM Identity Center

appear as `IdentityCenterUser` with full `onBehalfOf` details in the `userIdentity` element.

Discovering Amazon Quick events in your trail

Amazon Quick adds new event names as features launch and evolve. The non-API events listed below represent a subset of what your trail captures. To discover all Amazon Quick events in your account, query your trail for `eventSource: quicksight.amazonaws.com` and inspect the `eventName` values. Event names follow a consistent pattern: the operation name (such as `CreateFlow` or `UpdateAgent`) matches the feature it relates to.

Example: Amazon Quick log file entries

A trail is a configuration that enables delivery of events as log files to an Amazon S3 bucket that you specify. CloudTrail log files contain one or more log entries. An event represents a single request from any source and includes information about the requested action, the date and time of the action, request parameters, and so on. CloudTrail log files aren't an ordered stack trace of the public API calls, so they don't appear in any specific order.

The following example shows a CloudTrail log entry that demonstrates the `BatchCreateUser` action.

```
{
  "eventVersion": "1.11",
  "userIdentity": {
    "type": "Root",
    "principalId": "123456789012",
    "arn": "arn:aws:iam::123456789012:root",
    "accountId": "123456789012",
    "userName": "test-username"
  },
  "eventTime": "2017-04-19T03:16:13Z",
  "eventSource": "quicksight.amazonaws.com",
  "eventName": "BatchCreateUser",
  "awsRegion": "us-west-2",
  "requestParameters": null,
  "responseElements": null,
  "eventID": "e7d2382e-70a0-3fb7-9d41-a7a913422240",
  "readOnly": false,
  "eventType": "AwsServiceEvent",
```

```
"recipientAccountId":"123456789012",
"serviceEventDetails":
{
  "eventRequestDetails":
  {
    "users":
    {
      "test-user-11":
      {
        "role":"USER"
      },
      "test-user-22":
      {
        "role":"ADMIN"
      }
    }
  },
  "eventResponseDetails":
  {
    "validUsers":[
    ],
    "InvalidUsers":[
      "test-user-11",
      "test-user-22"
    ]
  }
}
```

Logging Amazon Quick data events in CloudTrail

In addition to management events, Amazon Quick supports CloudTrail data events for selected resource types. Data events are off by default. Enable them through advanced event selectors on your trail.

The following Amazon Quick resource types support data events:

Resource type	Event name	Availability
AWS::Quicksight::ActionConnector	InvokeAction	Generally available

Resource type	Event name	Availability
AWS::QuickSight::Flow	–	Generally available
AWS::QuickSight::FlowSession	–	Generally available
AWS::QuickSight::Companion	–	Varies by Region
AWS::QuickSight::CompanionAccess	–	Varies by Region
AWS::QuickSight::Namespaces	–	Varies by Region

Important

Data-event resource-type availability varies by account and AWS Region. If a resource type is not supported in your Region, the `put-event-selectors` call returns an error for that type. Build one `put-event-selectors` call that includes only the resource types supported in your Region, and omit unsupported types from that call.

Note

The casing of Quicksight compared with QuickSight in resource type ARNs is intentional and varies by resource type. Use the exact casing shown above when configuring event selectors.

Enabling data events

Use advanced event selectors to enable data events. Include a management-events selector to retain management event logging:

```
aws cloudtrail put-event-selectors \
```

```
--trail-name your-trail-name \
--region us-east-1 \
--advanced-event-selectors '[
  { "Name": "All management events",
    "FieldSelectors": [ { "Field": "eventCategory", "Equals": ["Management"] } ] },
  { "Name": "Quick connector data events",
    "FieldSelectors": [
      { "Field": "eventCategory", "Equals": ["Data"] },
      { "Field": "resources.type", "Equals":
["AWS::Quicksight::ActionConnector"] } ] },
  { "Name": "Quick flow data events",
    "FieldSelectors": [
      { "Field": "eventCategory", "Equals": ["Data"] },
      { "Field": "resources.type", "Equals": ["AWS::QuickSight::Flow"] } ] },
  { "Name": "Quick flow session data events",
    "FieldSelectors": [
      { "Field": "eventCategory", "Equals": ["Data"] },
      { "Field": "resources.type", "Equals": ["AWS::QuickSight::FlowSession"] } ] }
]'
```

Add selectors for `AWS::Quicksight::Companion`, `AWS::Quicksight::CompanionAccess`, and `AWS::QuickSight::Namespace` if supported in your Region.

InvokeAction data event

When a user or flow invokes a connector, CloudTrail records an `InvokeAction` data event. This event captures the following information:

- Caller identity, including `onBehalfOf.userId` and `endUserArn`
- The `actionConnectorId` and `actionId` of the invoked action
- The outcome: `invokeActionStatus` and `httpStatusCode`
- `readOnly`: `false` (always)
- `eventCategory`: `Data`, `managementEvent`: `false`
- Resource type: `AWS::Quicksight::ActionConnector`

⚠ Important

InvokeAction does not include the action's input parameters (such as an email recipient, subject, or body) or any approval decision. Use CloudWatch vended logs for conversation content.

ℹ Note

The `readOnly` field is always `false` for action connector invocations, including read-only actions such as listing emails. To distinguish read from write operations, filter by `actionId` rather than the `readOnly` field.

The following example shows an InvokeAction data event for a successful email action:

```
{
  "eventVersion": "1.11",
  "userIdentity": {
    "type": "IdentityCenterUser",
    "accountId": "111122223333",
    "onBehalfOf": {
      "userId": "a1b2c3d4-e5f6-7890-abcd-ef1234567890",
      "identityStoreArn": "arn:aws:identitystore::111122223333:identitystore/d-1234567890"
    }
  },
  "eventTime": "2026-08-19T20:26:36Z",
  "eventSource": "quicksight.amazonaws.com",
  "eventName": "InvokeAction",
  "awsRegion": "us-east-1",
  "sourceIPAddress": "203.0.113.1",
  "requestParameters": {
    "awsAccountId": "111122223333",
    "actionConnectorId": "52bde8f0-3de5-4e3d-9283-222f81a9fc4f",
    "actionId": "SendUserEmail",
    "endUserArn": "arn:aws:quicksight:us-east-1:111122223333:user/default/johndoe"
  },
  "responseElements": {
    "invokeActionStatus": "SUCCESS",
    "invokeActionOutput": {
```

```
        "httpInvokeActionOutput": {
            "httpStatusCode": "200"
        }
    },
    "readOnly": false,
    "resources": [
        {
            "accountId": "111122223333",
            "type": "AWS::Quicksight::ActionConnector",
            "ARN": "arn:aws:quicksight:us-east-1:111122223333:action-connector/52bde8f0-3de5-4e3d-9283-222f81a9fc4f"
        }
    ],
    "eventType": "AwsApiCall",
    "managementEvent": false,
    "eventCategory": "Data"
}
```

Tracking non-API events by using CloudTrail logs

The following non-API events are captured in your CloudTrail trail:

User management

- **CreateAccount** – Create Account
- **BatchCreateUser** – Create User
- **BatchResendUserInvite** – Invite User
- **UpdateGroups** – Update Groups

This event works with Enterprise edition only.

- **UpdateSpiceCapacity** – Update SPICE Capacity
- **DeleteUser** – Delete User
- **Unsubscribe** – Unsubscribe User

Subscription

- **CreateSubscription** – Create Subscription
- **UpdateSubscription** – Update Subscription
- **DeleteSubscription** – Delete Subscription

Dashboard

- **GetDashboard** – Get Dashboard
- **CreateDashboard** – Create Dashboard
- **UpdateDashboard** – Update Dashboard
- **UpdateDashboardAccess** – Update Dashboard Access
- **DeleteDashboard** – Delete Dashboard

Analysis

- **GetAnalysis** – Get Analysis
- **CreateAnalysis** – Create Analysis
- **UpdateAnalysisAccess** – Update Analysis Access
- **UpdateAnalysis** – Update Analysis
 - **RenameAnalysis** – Rename Analysis
 - **CreateVisual** – Create Visual
 - **RenameVisual** – Rename Visual
 - **DeleteVisual** – Delete Visual

- **DeleteAnalysis** – Delete Analysis

Data source

- **CreateDataSource** – Create Data Source
 - **FlatFile** – Flat file
 - **External** – External
 - **S3** – S3
 - **ImportS3ManifestFile** – S3 Manifest File
 - **Presto** – Presto
 - **RDS** – RDS
 - **Redshift** – Redshift (manual)
- **UpdateDataSource** – Update Data Source
- **DeleteDataSource** – Delete Data Source

Data set

- **CreateDataSet** – Create Data Set
 - **CustomSQL** – Custom SQL
 - **SQLTable** – SQL Table
 - **File** – CSV or XLSX
- **UpdateDataSet** – Update SQL Join Dataset

- **DeleteDataSet** – Delete Dataset
- **Querydatabase** – During a dataset refresh, query data source.

Monitoring Amazon Quick using CloudWatch metrics

Amazon Quick sends metrics to Amazon CloudWatch in the AWS/QuickSight namespace. These metrics cover dashboards, visuals, dataset ingestions, SPICE capacity, knowledge base indexing, connectors, and automations. Use them to observe availability and performance, and to create alarms for your Amazon Quick environment.

For more information about using Amazon CloudWatch, see the [Amazon CloudWatch User Guide](#).

Accessing Quick metrics in Amazon CloudWatch

Use the following procedure to access Amazon Quick metrics in Amazon CloudWatch.

To access Amazon Quick metrics in CloudWatch

1. Sign in to the AWS account that's associated with your Amazon Quick account.
2. In the upper-left corner of the AWS Management Console home page, choose **Services**, and then choose **CloudWatch**.
3. In the navigation pane, choose **Metrics**, **All metrics**, **QuickSight**.

Topics

- [Graph metrics with the Amazon CloudWatch console](#)
- [Creating alarms with the Amazon CloudWatch console](#)
- [Metrics](#)
- [Aggregate metrics](#)
- [Aggregate SPICE metrics](#)
- [Dimensions](#)

Graph metrics with the Amazon CloudWatch console

You can also use the Amazon CloudWatch console to graph metric data generated by Quick. For more information, see [Graphing metrics](#) in the *Amazon CloudWatch User Guide*.

Creating alarms with the Amazon CloudWatch console

You can create a Amazon CloudWatch alarm that monitors CloudWatch metrics for your Quick assets. When the metric reaches a threshold that you specify, CloudWatch automatically sends you a notification. For examples, see [Creating Amazon CloudWatch alarms](#) in the *Amazon CloudWatch User Guide*.

Metrics

The AWS/QuickSight namespace includes the following metrics for monitoring traffic and latency of your Amazon Quick assets.

Topics

- [Per-dashboard metrics](#)
- [Per-dataset ingestion metrics](#)
- [Per-visual metrics](#)
- [Per-unstructured dataset metrics](#)
- [Per-connector metrics](#)
- [Per-automation metrics](#)

Per-dashboard metrics

The following metrics track dashboard view counts and load times. You can find these metrics under the AWS/QuickSight/Dashboard Metrics group in CloudWatch.

Metric	Description	Dimension	Unit
DashboardViewCount	The number of times that a dashboard has been viewed. This number includes all access patterns such as web, mobile, and embedded. The most useful statistic for this metric is SUM, which	DashboardId	Count

Metric	Description	Dimension	Unit
	represents the total number of dashboard views during a set period of time.		
DashboardViewLoadTime	The amount of time that it takes for a Amazon Quick dashboard to load. The measurement begins when a user reaches the Amazon Quick dashboard and ends when all of the dashboard's visuals finish rendering. The most useful statistic for this metric is AVERAGE, which represents the average load time of a Quick dashboard during a set period of time.	DashboardId	Millisecond

Per-dataset ingestion metrics

The following metrics track ingestions for specific [SPICE](#) datasets. You can find these metrics under the AWS/QuickSight/Ingestion Metrics group in CloudWatch.

Metric	Description	Dimension	Unit
IngestionErrorCount	The number of failed ingestions.	DatasetId	Count

Metric	Description	Dimension	Unit
	The most useful statistic for this metric is SUM, which represents the total number of failed ingestions during a set period of time.		
IngestionInvocationCount	The number of ingestions that have been initiated. This includes scheduled and manual ingestions that are initiated through the console and the Amazon Quick API operations. The most useful statistic for this metric is SUM, which represents the total number of ingestions initiated during a set period of time.	DatasetId	Count

Metric	Description	Dimension	Unit
IngestionLatency	The time period between the initiation of an ingestion to the completion of the ingestion. The most useful statistic for this metric is AVERAGE, which represents the average runtime of ingestions during a set period of time.	DatasetId	Second
IngestionRowCount	The number of successful row ingestions. The most useful statistic for this metric is SUM, which represents the total amount of data ingested during a set period of time.	DatasetId	Count

Per-visual metrics

The following metrics track the load times and error counts of individual visuals on a Amazon Quick dashboard. You can find these metrics under the AWS/QuickSight/Visual Metrics group in CloudWatch.

Metric	Description	Dimension	Unit
VisualLoadTime	The time that it takes for a Amazon Quick visual to receive the necessary query data for an initial paint of the visual. This includes the round-trip query time from the client, to the Amazon Quick service, and then back to client. The most useful statistic for this metric is AVERAGE, which represents the average load time of a visual during a set period of time.	- DashboardId - SheetId - VisualId	Millisecond
VisualLoadErrorCount	The number of times that a Amazon Quick visual fails to complete a data query for the initial paint. Any error that occurs during a visual's loading period is included in this metric. The most useful statistic for this metric is SUM, which	- DashboardId - SheetId - VisualId	Count

Metric	Description	Dimension	Unit
	represents the total number of failed visual loads during a set period.		

Per-unstructured dataset metrics

The following metrics track document statistics and indexing status for Amazon Quick unstructured datasets. You can find these metrics under the AWS/QuickSight/QuickInstanceId group in CloudWatch.

Metric	Description	Dimension	Unit
QuickIndexDocumentCount	The number of documents in the unstructured Quick index.	- KnowledgeBaselId - QuickInstanceId	Count
QuickIndexExtractedTextSize	The extracted text size of the unstructured Quick index.	- KnowledgeBaselId - QuickInstanceId	Bytes
QuickIndexPurchaseInMB	The amount of storage that has been purchased for the unstructured Quick index.	- KnowledgeBaselId - QuickInstanceId	MB
QuickIndexCapacityConsumedRawFileSizeInGB	The raw file size consumed by the unstructured Quick index.	QuickInstanceId	GB
QuickIndexCapacityRawFileSizeLimitInGB	The raw file size limit of the unstructured Quick index.	QuickInstanceId	GB

Metric	Description	Dimension	Unit
DocumentsCrawled	The number of uploaded documents crawled in the unstructured Quick index.	- KnowledgeBaseld - QuickInstanceId	Count
DocumentsIndexed	The number of documents indexed in the unstructured Quick index.	- KnowledgeBaseld - QuickInstanceId	Count
DocumentsDeleted	The number of documents deleted from the unstructured Quick index.	- KnowledgeBaseld - QuickInstanceId	Count
DocumentsModified	The number of documents modified in the unstructured Quick index.	- KnowledgeBaseld - QuickInstanceId	Count
DocumentsFailedToIndex	The number of documents that failed to index in the unstructured Quick index.	- KnowledgeBaseld - QuickInstanceId	Count
ExtractedTextSize	The total text size extracted during a connector level sync in the unstructured Quick index.	- KnowledgeBaseld - QuickInstanceId	MB

Per-connector metrics

The following metrics track the number of invocations made to Quick connectors. You can find these metrics under the AWS/QuickSight namespace in CloudWatch.

Metric	Description	Dimension	Unit
ActionInvocationCount	The number of times a connector was invoked.	- ActionConnectorId - ActionConnectorType - ActionId	Count
ActionInvocationError	The number of times a connector failed to invoke.	- ActionConnectorId - ActionConnectorType - ActionId - InvokeErrorCode	Count

Per-automation metrics

The following metrics track automation run counts and duration. You can find these metrics under the AWS/QuickSight namespace in CloudWatch.

Metric	Description	Dimension	Unit
TotalAutomationRunCount	The total number of automation runs initiated. The most useful statistic for this metric is SUM, which represents the total number of automation runs during a set period of time.	AutomationId	Count

Metric	Description	Dimension	Unit
SuccessfulAutomationRunCount	The number of automation runs that completed successfully. The most useful statistic for this metric is SUM, which represents the total number of successful runs during a set period of time.	AutomationId	Count
FailedAutomationRunCount	The number of automation runs that failed. The most useful statistic for this metric is SUM, which represents the total number of failed runs during a set period of time.	AutomationId	Count

Metric	Description	Dimension	Unit
StoppedAutomationRunCount	The number of automation runs that were stopped before completion. The most useful statistic for this metric is SUM, which represents the total number of stopped runs during a set period of time.	AutomationId	Count
AutomationRunDuration	The elapsed time of an automation run from start to completion. The most useful statistic for this metric is AVERAGE, which represents the average duration of automation runs during a set period of time.	AutomationId	Seconds

Aggregate metrics

The AWS/QuickSight namespace includes the following aggregate metrics for monitoring traffic and latency of your Amazon Quick assets.

Topics

- [Aggregate dashboard metrics](#)
- [Aggregate ingestion metrics](#)

- [Aggregate visual metrics](#)
- [Aggregate unstructured dataset metrics](#)
- [Aggregate connector metrics](#)

Aggregate dashboard metrics

The following metrics track view counts and load times of all dashboards in a Amazon Quick account and region. You can find these metrics under the AWS/QuickSight/Aggregate Metrics group in CloudWatch.

Metric	Description	Unit
DashboardViewCount	The number of times that all Amazon Quick dashboards have been viewed across the entire Amazon Quick account in the region. This number is an aggregate that includes all access patterns such as web, mobile, and embedded. The most useful statistic for this metric is SUM, which represents the total number of Amazon Quick dashboard views during a set period of time.	Count
DashboardViewLoadTime	The amount of time that it takes for all Amazon Quick dashboards to load. The measurement begins when a user navigates to the Amazon Quick dashboard and ends when all of the dashboard's visuals finish rendering.	Millisecond

Metric	Description	Unit
	The most useful statistic for this metric is AVERAGE, which represents the average load time of all Amazon Quick dashboard during a set period of time.	

Aggregate ingestion metrics

The following metrics track all ingestions associated with a Amazon Quick account and AWS Region. You can find these metrics under the AWS/QuickSight/Aggregate Metrics group in CloudWatch.

Metric	Description	Unit
IngestionErrorCount	The number of failed ingestions. The most useful statistic for this metric is SUM, which represents the total number of failed ingestion during a set period.	Count
IngestionInvocationCount	The number of ingestion s that have been initiated. This includes scheduled and manual ingestions that are initiated through the console and the Amazon Quick API operations. The most useful statistic for this metric is SUM, which represents the total number	Count

Metric	Description	Unit
	of ingestions initiated during a set period of time.	
IngestionLatency	The time period between the initiation of an ingestion to the completion of the ingestion. The most useful statistic for this metric is AVERAGE, which represents the average runtime of ingestions during a set period of time.	Second
IngestionRowCount	The number of successful row ingestions. The most useful statistic for this metric is SUM, which represents the total amount of data ingested during a set period of time.	Count

Aggregate visual metrics

The following metrics track load times and error counts of all visuals on a dashboard and in a Amazon Quick account in a Region. You can find these metrics under the AWS/QuickSight/Aggregate Metrics group for CloudWatch.

Metric	Description	Unit
VisualLoadTime	The time that it takes for all Amazon Quick visuals to receive the necessary query data for an initial paint of the visuals. This includes the	Millisecond

Metric	Description	Unit
	round-trip query time from the client, to the Amazon Quick service, and then back to the client. The most useful statistic for this metric is AVERAGE, which represents the average load time of all visuals during a set period of time.	
VisualLoadErrorCount	The number of times that all Amazon Quick visuals that belong to the Amazon Quick account fail to complete a data query for an initial paint. The most useful statistic for this metric is SUM, which represents the total number of failed visuals during a set period.	Count

Aggregate unstructured dataset metrics

The following metrics track all unstructured dataset metrics within a Amazon Quick account in a Region. You can find these metrics under the AWS/QuickSight/Aggregate Metrics group for CloudWatch.

Metric	Description	Unit
QuickIndexDocumentCount	The number of documents in the unstructured Quick index. The most useful statistic for this metric is SUM, which	Count

Metric	Description	Unit
	represents the total number of documents added to your index during a set period of time.	
QuickIndexExtractedTextSize	The extracted text size of the unstructured Quick index. The most useful statistic for this metric is SUM, which represents the total size of all text across all documents in your index.	Bytes
QuickIndexPurchasedInMB	The amount of storage that has been purchased for the unstructured Quick index. The most useful statistic for this metric is SUM, which represents the total size of purchased storage in MB across your index.	MB
QuickIndexCapacity ConsumedRawFileSizeInGB	The raw file size consumed by the unstructured Quick index. The most useful statistic for this metric is SUM, which represents the total raw file size consumed across your index in GB.	GB

Metric	Description	Unit
QuickIndexCapacityRawFileSizeLimitInGB	The raw file size limit of the unstructured Quick index. The most useful statistic for this metric is SUM, which represents the total raw file size limit across your index in GB.	GB
DocumentsCrawled	The number of uploaded documents crawled in the unstructured Quick index. The most useful statistic for this metric is SUM, which represents the total number of documents crawled in your index.	Count
DocumentsIndexed	The number of documents indexed in the unstructured Quick index. The most useful statistic for this metric is SUM, which represents the total number of documents indexed.	Count

Metric	Description	Unit
DocumentsDeleted	The number of documents deleted from the unstructured Quick index. The most useful statistic for this metric is SUM, which represents the total number of documents deleted from your index.	Count
DocumentsModified	The number of documents modified in the unstructured Quick index. The most useful statistic for this metric is SUM, which represents the total number of documents modified in your index.	Count
DocumentsFailedToIndex	The number of documents that failed to index in the unstructured Quick index. The most useful statistic for this metric is SUM, which represents the total number of documents that failed to index.	Count

Metric	Description	Unit
ExtractedTextSize	The total text size extracted during a connector level sync in the unstructured Quick index. The most useful statistic for this metric is SUM, which represents the total size of documents extracted during a connector level sync.	MB

Aggregate connector metrics

The following metrics track all Quick connector invocations associated with a Amazon Quick account in an AWS Region. You can find these metrics under the AWS/QuickSight/Aggregate Metrics group for CloudWatch.

Metric	Description	Unit
ActionInvocationCount	The number of connector invocations made. The most useful statistic for this metric is SUM, which represents the total number of connector invocations initiated during a set period of time.	Count
ActionInvocationError	The number of failed connector invocations. The most useful statistic for this metric is SUM, which represents the total number of connector invocations that	Count

Metric	Description	Unit
	failed during a set period of time.	

Aggregate SPICE metrics

Use the following metrics to monitor SPICE consumption and avoid reaching the consumption limit that can cause ingestions to fail. CloudWatch stores statistics for up to 15 months, so you can access historical information to better understand the consumption trends of your Amazon Quick account. You can find these metrics in the `AWS/QuickSight/Aggregate Metrics` group for CloudWatch.

Metric	Description	Unit
SPICECapacityLimitInMB	This value represents the provisioned SPICE capacity at a specific point in time. This metric refreshes when an update of 1 MB or more in consumed or purchased capacity is made.	MegaBytes
SPICECapacityConsumedInMB	This value represents the consumed SPICE capacity at a specific point in time. This metric refreshes when an update of 1 MB or more in consumed or purchased capacity is made.	MegaBytes

Dimensions

The following table lists Quick metric dimensions that appear in Amazon CloudWatch.

Dimension	Description
DashboardId	The public ID of a Quick dashboard. You can use the <code>ListDashboards</code> API operation to see a list of every dashboard in your Amazon Quick account. For more information, see ListDashboards in the <i>Amazon Quick API Reference</i>.
DatasetId	The public ID of a Quick dataset. You can use the <code>ListDataSets</code> API operation to see a list of every dataset in your Amazon Quick account. For more information, see ListDataSets in the <i>Amazon Quick API Reference</i>.
SheetId	The public ID of a Quick sheet.
VisualId	The public ID of a Quick visual.
KnowledgeBaseId	The public ID of a Quick knowledge base.
QuickInstanceId	The public ID of the Quick instance.
ActionConnectorId	The public ID of the Quick connector.
ActionConnectorType	The type of the Quick connector.
ActionId	The public ID of the Quick Action.
InvokeErrorCode	The error code related to a failed Quick Action Connector invocation.

Compliance validation for Amazon Quick

Third-party auditors assess the security and compliance of Amazon Quick as part of multiple AWS compliance programs. Amazon Quick, across its web, desktop, and mobile clients, is in scope for the following programs:

- Federal Risk and Authorization Management Program (FedRAMP)
- Health Insurance Portability and Accountability Act (HIPAA)
- Payment Card Industry Data Security Standard (PCI DSS)
- System and Organization Controls (SOC) 1, SOC 2, and SOC 3
- International Organization for Standardization (ISO) 9001, ISO 27001, ISO 27018, and ISO 27019
- Cloud Computing Compliance Criteria Catalogue (C5) from the German Federal Office for Information Security (BSI)
- Health Information Trust Alliance Common Security Framework (HITRUST CSF)

When a program lists Quick as in scope, the audits and assessments for that program include the service. In-scope status does not make your own workloads compliant. Under the [shared responsibility model](#), you configure Quick and manage your data to meet the requirements that apply to you.

Compliance programs that include Quick

The following programs include Quick. For each program, use the linked AWS page to confirm the current status, the authorized AWS Regions, and any program details before you rely on the program for your own workloads.

FedRAMP

Quick is in scope for the FedRAMP program. To confirm the current authorization status and the authorized AWS Regions, see [AWS services in scope for FedRAMP](#). You remain responsible for meeting the FedRAMP requirements that apply to your workloads.

HIPAA eligibility

Quick is a HIPAA Eligible Service. This eligibility applies across the Quick web, desktop, and mobile clients, subject to the shared responsibility model and your configuration. To process, store, or transmit protected health information (PHI), you must have an AWS Business Associate

Addendum (BAA) in place. You must also configure the service according to your obligations. To confirm current eligibility, see the [HIPAA Eligible Services Reference](#). For more information about using AWS to build HIPAA workloads, see [HIPAA Overview](#).

PCI DSS

Quick is in scope for the Payment Card Industry Data Security Standard (PCI DSS). To confirm the current status, see [AWS services in scope for PCI DSS](#). You are responsible for the security of the cardholder data that you handle in the service.

SOC reports

Quick is in scope for the AWS System and Organization Controls (SOC) reports, including SOC 1, SOC 2, and SOC 3. To confirm the current status and download the reports that describe the AWS controls, see [AWS services in scope for SOC](#).

ISO 9001, ISO 27001, ISO 27018, and ISO 27019

Quick is in scope for the ISO 9001, ISO 27001, ISO 27018, and ISO 27019 standards. ISO 27018 is the code of practice for protecting personal data in the cloud. To confirm the current status and review the certifications, see [ISO 27001 Overview](#) and the [AWS services in scope by compliance program](#) page.

BSI C5

Quick is in scope for the AWS Cloud Computing Compliance Criteria Catalogue (BSI C5). To confirm the current status, see [AWS services in scope for BSI C5](#).

HITRUST CSF

Quick is in scope for the HITRUST Common Security Framework (CSF). To confirm the current status, see [AWS services in scope for HITRUST CSF](#).

Verifying current reports and scope

Program scope and audit reports change over time. To confirm the status that applies to you, use the following AWS resources.

- To download third-party audit reports, use AWS Artifact. For more information, see [Downloading reports in AWS Artifact](#).
- For the most current list of AWS services in scope of specific compliance programs, see [AWS services in scope by compliance program](#). For general information, see [AWS compliance programs](#).

Your compliance responsibilities

Your compliance responsibility when you use Amazon Quick depends on the sensitivity of your data, your company's compliance objectives, and applicable laws and regulations. Use the following AWS resources to help you meet your obligations:

- [Security and compliance quick start guides](#) – These deployment guides discuss architectural considerations and provide steps for deploying security- and compliance-focused baseline environments on AWS.
- [AWS compliance resources](#) – This collection of workbooks and guides might apply to your industry and location.
- [AWS Config](#) – This AWS service assesses how well your resource configurations comply with internal practices, industry guidelines, and regulations.
- [AWS Security Hub CSPM](#) – This AWS service provides a comprehensive view of your security state within AWS that helps you check your compliance with security industry standards and best practices.

Resilience in Amazon Quick

Quick is built by AWS and runs on AWS-managed infrastructure. It takes full advantage of the high availability features provided by AWS.

The AWS global infrastructure is built around AWS Regions and Availability Zones. AWS Regions provide multiple physically separated and isolated Availability Zones, which are connected with low-latency, high-throughput, and highly redundant networking. With Availability Zones, you can design and operate applications and databases that automatically fail over between Availability Zones without interruption. Availability Zones are more highly available, fault tolerant, and scalable than traditional single or multiple data center infrastructures.

Because Amazon Quick is an AWS-managed application, all patches and updates are applied by AWS as needed.

For more information about AWS Regions and Availability Zones, see [AWS global infrastructure](#).

Infrastructure security in Amazon Quick

Intended audience: Amazon Quick administrators

Quick is delivered as a web application, hosted on dedicated Amazon EC2 hosts, separate from AWS virtual private clouds (VPCs). Instead of deploying Amazon Quick on your own hosts, you access the Amazon Quick service through Regional public endpoints. Amazon Quick accesses data sources over a secured internet connection from Regional endpoints. To access data sources that are located inside a corporate network, configure the network to allow access from one of the Amazon Quick public IP address blocks. We recommend that you consider using a VPC (a virtual network dedicated to your AWS account).

For more information, see the following:

- [Global Infrastructure: The Most Extensive, Reliable, and Secure Global Cloud Infrastructure](#)
- [AWS Regions, websites, IP address ranges, and endpoints](#)
- [Connecting to a Amazon VPC with Amazon Quick](#)

As a managed service, Quick is protected by the AWS global network security procedures that are described in the [Amazon Web Services: Overview of Security Processes](#) paper.

If you use AWS published API calls to access Amazon Quick through the network, clients must support Transport Layer Security (TLS) 1.2 or later. Clients must also support cipher suites with perfect forward secrecy (PFS) such as Ephemeral Diffie-Hellman (DHE) or Elliptic Curve Ephemeral Diffie-Hellman (ECDHE). Most modern systems such as Java 7 and later support these modes.

Additionally, requests must be signed by using an access key ID and a secret access key that is associated with an AWS Identity and Access Management (IAM) principal. Or you can use the [AWS Security Token Service](#) (AWS STS) to generate temporary security credentials to sign requests.

You can call these API operations from any network location, but Amazon Quick does support resource-based access policies, which can include restrictions based on the source IP address. You can also use Amazon Quick policies to control access from specific Amazon Virtual Private Cloud (Amazon VPC) endpoints or specific VPCs. Effectively, this isolates network access to a given Amazon Quick resource from only the specific VPC within the AWS network. For more information on using Amazon Quick in a VPC, see [Connecting to a Amazon VPC with Amazon Quick](#).

Security in Quick

Amazon Quick provides a number of security features to consider as you develop and implement your own security policies. The following best practices are general guidelines and don't represent a complete security solution. Because these best practices might not be appropriate or sufficient for your environment, treat them as helpful considerations rather than prescriptions.

For more information about how your data is protected, including the statement that customer data is not used for training or improving underlying LLMs, see [Data protection in Amazon Quick](#).

For a conceptual explanation of how AI and agent capabilities change the risk analysis for Quick, including the two-condition model for adversarial content risk, see [AI and agent security in Amazon Quick](#).

Decide before you deploy – The following decisions are irreversible or difficult to change after initial configuration. Make them deliberately before you create resources or onboard users.

Register encryption keys before you create resources – Quick encrypts all data at rest by default. For more information, see [Encryption at rest](#). The following decision is about whether you or AWS controls the keys. If you require control of the encryption keys that protect your data, register a customer managed KMS key as the account default *before* you create resources. Two key paths matter:

- **Account default key** – Participating resource types select the account default key when data is written. Changing the default affects future writes but does not re-encrypt existing content.
- **Amazon Q data key** – Set once, the first time Amazon Q data is created in the account, from the default key registered at that moment. You cannot change this key afterward. For more information, see [Amazon Q data key](#).
- **Research resources** – A Research resource associates a key at creation. A Research resource created without a key association cannot be moved to one later.

Customer managed KMS key support requires an Enterprise subscription. For the full list of participating resource types, see [Customer managed KMS key scope](#).

Choose your identity model – Choose deliberately among IAM Identity Center, Active Directory, IAM federation, and Quick local users. Provisioning and offboarding behavior differ by model. In particular, Quick detects inactive users and populates the **Inactive users** list. This detection applies only to accounts that federate through IAM Identity Center or Active Directory. For more information, see [User lifecycle and data handling in Amazon Quick](#).

Select the knowledge base access model at creation – For supported integrations, decide at creation whether a knowledge base is access control list (ACL) aware. Document-level ACL configuration is permanent. You cannot turn on ACLs for a knowledge base created without ACL support, and you cannot turn them off after you turn them on. To change ACL configuration, create a new knowledge base with your desired setting from the start.

- **Supported integrations** – Amazon S3, Google Drive, Microsoft SharePoint, Atlassian Confluence Cloud, and Microsoft OneDrive.
- **Sharing and retrieval are separate controls** – Knowledge base sharing determines who can use the knowledge base. For ACL-aware knowledge bases, source-document ACLs further limit which indexed documents each authorized user can retrieve. Review both controls before granting access.
- **Fail-closed behavior** – If Quick cannot evaluate document permissions for a query, it returns no documents rather than unfiltered results.
- **Sync cadence** – Quick synchronizes identity and document-permission changes on the knowledge base refresh schedule, which is every 24 hours by default. Configure a shorter schedule when your access-change requirements demand it.
- **Shared email addresses** – If multiple Quick users share the same email address within a namespace, the system denies access to everyone using that shared email.
- **ACL resolution scope** – Quick resolves all ACLs within the namespace of the knowledge base creator.
- **Email recycling** – If your organization reassigns an email address before the next ACL refresh and the previous holder never used Quick for chat, temporary access leakage is possible. Update ACLs and refresh the knowledge base before reassigning the email.
- **Research incompatibility** – Knowledge bases with document-level ACLs enabled are not compatible with Quick Research.

For more information, see [Best practices for managing ACLs in knowledge bases](#).

Restrict access – After your account is configured, restrict access to the minimum required for each role.

Apply least privilege – Assign Quick roles that match each user's responsibilities. Use custom permissions profiles to further scope down the default capabilities available to a role. Custom permissions can only restrict. They cannot grant capabilities a user does not already have. When you assign profiles at multiple levels, the most specific level wins: user-level overrides role-level, which overrides account-level. Custom permissions require an Enterprise subscription. For more information, see [Custom permissions](#).

For BI workloads, supplement permissions with row-level security, column-level security, and run-as roles to limit data access at query time.

Deny by default for AI capabilities – If your organization has model-risk or change-approval requirements, restrict the AI capability category so that new AI capabilities are denied on launch day until you explicitly allow them. Without this restriction, every new AI capability that Quick ships is automatically available to all users upon release. For more information, see [Deny by Default](#).

Configure AI guardrails – Built-in content safety screening applies automatically to chat interactions. You can further configure up to 50 blocked words and phrases that Quick checks against both user requests and generated responses. For more information, see [AI guardrails in Amazon Quick](#).

Match oversight to how the action runs – Quick supports two invocation models with different security characteristics. Match your review depth to the autonomy and consequence of each model.

- **On-demand actions** – User-initiated, interactive, and authenticated with the individual user's personal credentials. Each action runs as the invoking user and is bounded by that user's connector permissions. For more information, see [On-demand actions](#).
- **Automated workflows** – System-level, scheduled or event-triggered, non-interactive, and authenticated at service level. An automated workflow authenticates at service level. It does not act as an individual user, is not bounded by that user's connector permissions, and runs unattended. Give automated workflows closer scrutiny during review. For more information, see [Automated workflows](#).
- **Restricting autonomous execution** – If unattended execution is not acceptable for some or all users, use the following custom permissions capabilities to restrict it:
 - **Automate** – Prevents restricted users from creating, updating, or running automations.
 - **Flows** – Restricts creating, updating, sharing, and running flows.
 - **Triggers** – Restricts all trigger capabilities, including inbound email triggers and Quick event triggers.

Apply these restrictions in an account-level profile to restrict unattended execution across the entire account. For more information, see [Custom permissions](#).

- **Adding a review step** – For workflows where a human review step is appropriate, use the human-in-the-loop task center to pause execution and collect approval before the workflow continues. For more information, see [Human-in-the-loop task center](#).

Treat retrieved and external content as untrusted – The risk of acting on adversarial content requires two conditions: the assistant can read content from outside your trust boundary, and it can take a consequential action. Narrowing either side reduces exposure.

- **Read surface** – Content from the following sources enters the assistant's context: knowledge base and index retrieval, files uploaded to spaces and chat, web search, third-party research providers, and content returned by connectors.
- **Action surface** – Limit the connectors, actions, and outbound paths available to each role. A narrower action scope reduces the consequence of any single compromised input.
- **Built-in screening** – Built-in safety checks screen user requests, including requests that contain prompt attacks. Pair this screening with narrow action scope rather than relying on screening alone. For more information, see [AI guardrails in Amazon Quick](#).

Restrict network access – Use IP and VPC endpoint rules to control which networks can reach Quick through web, embedded, and mobile access. Rule types include CIDR blocks (IPv4), VPC IDs, and VPC endpoint IDs. You can add up to 100 IP and VPC endpoint rules combined. These rules are global and apply to all AWS Regions. Rule changes can take up to 10 minutes to take effect. IP and VPC endpoint rules require an Enterprise subscription.

IP and VPC endpoint rules do *not* restrict access to the public API. To restrict API access, use IAM policy conditions. For more information, see [Turning on Internet Protocol \(IP\) and VPC endpoint restrictions in Amazon Quick](#).

Transfer ownership before you offboard – Assets left in the account without an owner become inaccessible to every user. Transfer ownership of connectors, agents, and automations before you remove a user from the account.

- **Review orphaned assets** – Periodically review connectors, agents, and automations whose owners have left your organization. For more information, see [Orphaned assets](#).
- **Review dormant connections** – Managed OAuth refresh tokens carry a 90-day lifecycle. Use this cadence as a trigger to review whether each connection is still needed, and remove connections that are no longer in use.
- **Plan for the full lifecycle** – For detailed guidance on user removal and its effects, see [User lifecycle and data handling in Amazon Quick](#).

Review what leaves – Each outbound path is customer-enabled, and each connection has its own authentication and data-access characteristics. Review each path as a potential egress channel before you turn it on.

Require authentication on every connection – The authentication options available depend on the connector type. Some connector types allow connecting without authentication, including remote Model Context Protocol (MCP) servers, agent-to-agent connectors, and generic HTTP connectors. Confirm that every connection requires authentication before you enable it.

- **User-level OAuth** – Each user authenticates individually and acts as themselves. Prefer user-level authentication where the connector supports it, because the impact of a misconfigured connection is limited to a single user's access scope.
- **Shared client-credentials connection** – A single set of credentials is shared across all users of the connector. Every user of that connector inherits its reach and access scope. Scope shared credentials narrowly and review what they grant.
- **Per-connector restriction** – You can restrict individual connectors for Create and Update, Share, and Use through custom permissions, so a connector can be enabled for a narrow set of users rather than everyone. For more information, see [Custom permissions](#).

For more information about the available authentication methods, see [Authentication methods](#).

Review each outbound path before enabling it – Before you enable an outbound capability, verify that the destination and the data that can flow to it meet your organization's requirements. Outbound paths include connectors, remote MCP servers, web search, third-party research providers, public apps, embedding, and export destinations. For per-path detail on what data can leave, what triggers it, who controls it, and what to verify, see [Outbound connections and data egress in Amazon Quick](#).

Configure observability – Configure monitoring and logging before you onboard users at scale. Use the following signals together when you design monitoring and incident-response procedures. Before you rely on a signal for compliance, detection, or investigation, confirm that it covers the event you intend to track.

- **CloudTrail** – Records supported API operations and a documented set of non-API events such as dashboard views and user-management actions. CloudTrail does not capture chat conversation content or user feedback. For those signals, use CloudWatch vended logs. For more information, see [Monitoring Amazon Quick using CloudTrail](#).

- **CloudWatch vended logs** – Five log types deliver conversational and operational data to destinations you control (Amazon CloudWatch Logs, Amazon S3, or Firehose):
 - Chat conversations (CHAT_LOGS)
 - User feedback (FEEDBACK_LOGS)
 - Agent and research hours (AGENT_HOURS_LOGS)
 - Index storage usage (INDEX_USAGE_LOGS)
 - Knowledge base file sync (KB_FILE_SYNC_LOGS)

Vended logs do not flow by default. Configure delivery shortly after you enable Quick AI features so that you do not lose early interaction data. For more information, see [Monitoring Amazon Quick using CloudWatch Logs](#).

- **CloudWatch metrics and alarms** – Near-real-time operational metrics cover dashboards, visuals, dataset ingestions, unstructured datasets, connectors, and SPICE capacity. You can create CloudWatch alarms that send notifications when a metric reaches a threshold that you specify. For more information, see [Monitoring Amazon Quick using CloudWatch metrics](#).
- **Analytics dashboard** – Provides usage, adoption, and feedback insights for IAM administrators. Access requires the `quicksight:QuickSuiteUsageMetrics` permission. The analytics dashboard can expose detailed query text, failed-query content, and the full conversation context associated with negative feedback. Grant analytics access only to authorized administrators. For more information, see [Using the Amazon Quick analytics dashboard](#).
- **Feature-specific reports** – Some features produce their own operational reports, such as knowledge base sync reports that document per-file sync status.

You control retention for all externalized signals. Set destination retention policies to match your compliance requirements. CloudTrail event history retains 90 days of management events by default. CloudWatch Logs, Amazon S3, and Firehose destination retention follows the policies you configure at each destination. For more information, see [Incident response, logging, and monitoring in Amazon Quick](#).

CloudTrail records metadata (who, when, what operation) but not the message body or response content. CloudWatch vended chat logs capture user messages and system responses but not source document content. The analytics dashboard exposes conversation context only in the scope of negative feedback review. If your compliance model requires a signal not listed here, verify coverage before assuming it exists.

Data source connectivity – The following practices help you connect securely to data sources.

Use VPC connections for private sources – Use a virtual private cloud (VPC) connection for data in AWS data sources and for data in on-premises servers without public connectivity. For AWS sources, Quick uses an elastic network interface for secure, private communication with data sources in a VPC. For on-premises sources, use Direct Connect or AWS Site-to-Site VPN to create a secure, private link. VPC connections require an Enterprise subscription. For more information, see [Managing VPC connection in Amazon Quick](#).

Encrypt database connections with SSL – Use SSL to connect to your databases, especially if you are using public networks. Using SSL with Quick requires certificates signed by a publicly recognized certificate authority (CA).

Configure firewall rules for user and data access – To allow users to access Quick, allow access to HTTPS and WebSockets Secure (wss://) protocol. To allow Quick to reach a database on a non-AWS server, change that server's firewall configuration to accept traffic from the applicable Quick IP address range. For more information, see [Network and database configuration requirements](#).

AWS managed policies for Amazon Quick

To add permissions to users, groups, and roles, it is easier to use AWS managed policies than to write policies yourself. It takes time and expertise to [create IAM customer managed policies](#) that provide your team with only the permissions they need. To get started quickly, you can use our AWS managed policies. These policies cover common use cases and are available in your AWS account. For more information about AWS managed policies, see [AWS managed policies](#) in the *IAM User Guide*.

AWS services maintain and update AWS managed policies. You can't change the permissions in AWS managed policies. Services occasionally add additional permissions to an AWS managed policy to support new features. This type of update affects all identities (users, groups, and roles) where the policy is attached. Services are most likely to update an AWS managed policy when a new feature is launched or when new operations become available. Services do not remove permissions from an AWS managed policy, so policy updates won't break your existing permissions.

Additionally, AWS supports managed policies for job functions that span multiple services. For example, the **ReadOnlyAccess** AWS managed policy provides read-only access to all AWS services and resources. When a service launches a new feature, AWS adds read-only permissions for new operations and resources. For a list and descriptions of job function policies, see [AWS managed policies for job functions](#) in the *IAM User Guide*.

Topics

- [AWS managed policy: AWSQuickSightElasticsearchPolicy](#)
- [AWS managed policy: AWSQuickSightOpenSearchPolicy](#)
- [AWS managed policy: AWSQuickSightSageMakerPolicy](#)
- [AWS managed policy: AWSQuickSightAssetBundleExportPolicy](#)
- [AWS managed policy: AWSQuickSightAssetBundleImportPolicy](#)
- [Amazon Quick updates to AWS managed policies](#)

AWS managed policy: AWSQuickSightElasticsearchPolicy

This information is provided for backward compatibility only. The AWSQuickSightOpenSearchPolicy AWS managed policy replaces the AWSQuickSightElasticsearchPolicy AWS managed policy.

Previously, you used the AWSQuickSightElasticsearchPolicy AWS managed policy to provide access to Amazon Elasticsearch Service resources from Amazon Quick. Starting on or after September 7, 2021, Amazon Elasticsearch Service is renamed to Amazon OpenSearch Service.

Wherever you are using AWSQuickSightElasticsearchPolicy, you can update to the new AWS managed policy that's called AWSQuickSightOpenSearchPolicy. You can attach the policy to your IAM entities. Amazon Quick also attaches the policy to a service role that allows Amazon Quick to perform actions on your behalf. AWSQuickSightElasticsearchPolicy is still available and as of August 31, 2021, had the same permissions as the new policy. However, AWSQuickSightElasticsearchPolicy is no longer kept up-to-date with latest changes.

This policy grants read-only permissions that allow access to OpenSearch (previously known as Elasticsearch) resources from Amazon Quick.

Permissions details

This policy includes the following permissions:

- **es** – Allows principals to use `es:ESHttpGet` to access your OpenSearch (previously known as Elasticsearch) domains, cluster settings, and indices. This is required to use the search service from Amazon Quick.

- `es` – Allows principals to use `es:ListDomainNames` to list your OpenSearch (previously known as Elasticsearch) domains. This is required to initiate access of the search service from Amazon Quick.
- `es` – Allows principals to use `es:DescribeElasticsearchDomain` to search your OpenSearch (previously known as Elasticsearch) domains. This is required to use the search service from Amazon Quick.
- `es` – Allows principals to use `es:ESHttpPost` and `es:ESHttpGet` with your OpenSearch (previously known as Elasticsearch) domains. This is required to use a SQL plugin with read-only access to the search service domains from Amazon Quick.

For information on the contents of this IAM policy, see [AWSQuickSightElasticsearchPolicy](#) in the IAM console.

AWS managed policy: AWSQuickSightOpenSearchPolicy

Use the `AWSQuickSightOpenSearchPolicy` AWS managed policy to provide access to Amazon OpenSearch Service resources from Amazon Quick. `AWSQuickSightOpenSearchPolicy` replaces `AWSQuickSightElasticsearchPolicy`. As of August 31, 2021, this policy had the same permissions as the legacy policy, `AWSQuickSightElasticsearchPolicy`. For now, you can use them interchangeably. For the long term, we recommend updating your policy usage to `AWSQuickSightOpenSearchPolicy`.

You can attach `AWSQuickSightOpenSearchPolicy` to your IAM entities. Amazon Quick also attaches this policy to a service role that allows Amazon Quick to perform actions on your behalf.

This policy grants read-only permissions that allow access to OpenSearch resources from Amazon Quick.

Permissions details

This policy includes the following permissions:

- `es` – Allows principals to use `es:ESHttpGet` to access your OpenSearch domains, cluster settings, and indices. This is required to use Amazon OpenSearch Service from Amazon Quick.
- `es` – Allows principals to use `es:ListDomainNames` to list your OpenSearch domains. This is required to initiate access of Amazon OpenSearch Service from Amazon Quick.

- `es` – Allows principals to use `es:DescribeElasticsearchDomain` and `es:DescribeDomain` to search your OpenSearch domains. This is required to use Amazon OpenSearch Service from Amazon Quick.
- `es` – Allows principals to use `es:ESHttpPost` and `es:ESHttpGet` with your OpenSearch domains. This is required to use a SQL plugin with read-only access to Amazon OpenSearch Service domains from Amazon Quick.

For information on the contents of this IAM policy, see [AWSQuickSightOpenSearchPolicy](#) in the IAM console.

AWS managed policy: AWSQuickSightSageMakerPolicy

Use the `AWSQuickSightSageMakerPolicy` AWS managed policy to provide access to Amazon SageMaker AI resources from Amazon Quick.

You can attach `AWSQuickSightSageMakerPolicy` to your IAM entities. Amazon Quick also attaches this policy to a service role that allows Amazon Quick to perform actions on your behalf.

This policy grants read-only permissions that allow access to Amazon SageMaker AI resources from Amazon Quick.

To view the `AWSQuickSightSageMakerPolicy`, see [AWSQuickSightSageMakerPolicy](#) in the [AWS Managed Policy reference](#).

Permissions details

This policy includes the following permissions:

- `sagemaker` – .
- `s3` – Allows principals to use `s3:GetObject` on all Amazon S3 buckets that start with the prefix `arn:aws:s3:::sagemaker.*` to access data stored in SageMaker AI default buckets. This is required to load models shared from Amazon SageMaker AI Canvas to the default Amazon SageMaker AI Canvas Amazon S3 bucket.
- `s3` – Allows principals to use `s3:PutObject` to export objects into an Amazon S3 bucket. This is required to support existing datasets from Amazon Quick to Amazon SageMaker AI Canvas to build predictive models.
- `s3` – Allows principals to use `s3:ListBucket` to allow Amazon Quick to validate an existing Amazon SageMaker AI Canvas bucket in Amazon S3. This is required to allow the export of data from Amazon Quick to Amazon SageMaker AI Canvas to build predictive models.

- `s3` – Allows principals to use `s3:GetObject` on all Amazon Quick– owned Amazon S3 buckets that start with the prefix `arn:aws:s3:::quicksight-ml`. This is required to allow Amazon Quick to access the predictions that are generated by Amazon SageMaker AI Canvas. The generated predictions can be appended to a Amazon Quick dataset.
- `sagemaker` – Allows principals to use `sagemaker:CreateTransformJob`, `sagemaker:DescribeTransformJob`, and `sagemaker:StopTransformJob` to perform SageMaker AI transform jobs on your behalf. This is required for Amazon Quick to request predictions from SageMaker AI models that can be appended to a Amazon Quick dataset.
- `sagemaker` – Allows principals to use `sagemaker:ListModels` to list your SageMaker AI models. This is required to allow generated SageMaker AI models to appear in Amazon Quick.

AWS managed policy: AWSQuickSightAssetBundleExportPolicy

Use the `AWSQuickSightAssetBundleExportPolicy` AWS managed policy to perform asset bundle export operations. You can attach `AWSQuickSightAssetBundleExportPolicy` to your IAM entities.

This policy grants read-only permissions that allow access to Amazon Quick asset resources. To view the details of this policy, see [AWSQuickSightAssetBundleExportPolicy](#) in the AWS Managed Policy reference.

This policy includes the following permissions:

- `quicksight` – Allows principals to use `quicksight:Describe*` and `quicksight:List*` to find and fetch Amazon Quick assets and their corresponding permissions.
- `quicksight` – Allows principals to use `quicksight:ListTagsForResource` to fetch tags of Amazon Quick assets.
- `quicksight` – Allows principals to list, execute, and get the status of an Asset bundle export job. This policy uses the `quicksight:ListAssetBundleExportJob`, `StartAssetBundleExportJob`, and `quicksight:DescribeAssetBundleExportJob` permissions.

AWS managed policy: AWSQuickSightAssetBundleImportPolicy

Use the `AWSQuickSightAssetBundleImportPolicy` AWS managed policy to perform asset bundle import operations. This managed policy does not grant permissions for any run-as-role functionality with the `iam:passrole` that is required for some VPC connection and `DataSource`

operations. This policy also does not grant access to retrieve objects from a users Amazon S3 bucket.

You can attach the `AWSQuickSightAssetBundleImportPolicy` to your IAM entities. This policy grants read and write permissions that allow access to Amazon Quick resources. To view the details of this policy, see [AWSQuickSightAssetBundleImportPolicy](#) in the AWS Managed Policy reference.

This policy includes the following permissions:

- `quicksight` – Allows principals to use `quicksight:Describe*` and `quicksight:List*` to detect changes in the Amazon Quick assets and their permissions.
- `quicksight` – Allows principals to use `quicksight:Create*` and `quicksight:Update*` to make changes to the Amazon Quick assets and permissions from the supplied asset bundle.
- `quicksight` – Allows principals to use `quicksight:ListTagsForResource`, `quicksight:TagResource`, and `quicksight:UntagResource` to update the tags of Amazon Quick assets.
- `quicksight` – Allows principals to list, execute, and get the status of an Asset bundle import job. This policy uses the `quicksight:ListAssetBundleImportJob`, `quicksight:StartAssetBundleImportJob`, and `quicksight:DescribeAssetBundleImportJob` permissions.

Amazon Quick updates to AWS managed policies

View details about updates to AWS managed policies for Amazon Quick since this service began tracking these changes. For automatic alerts about changes to this page, subscribe to the RSS feed on the [Amazon Quick Document History](#) page.

Change	Description	Date
AWSQuickSightAssetBundleExportPolicy – New policy	Amazon Quick added new permissions to simplify Asset bundle export operations.	March 27, 2024
AWSQuickSightAssetBundleImportPolicy – New policy	Amazon Quick added new permissions to simplify Asset bundle import operations.	March 27, 2024

Change	Description	Date
AWSQuickSageMakerPolicy – Update to an existing policy	Amazon Quick added new permissions to allow integration with Amazon SageMaker AI Canvas.	July 25, 2023
AWSQuickSightElasticsearchPolicy – Update to an existing policy	Amazon Quick added new permissions to provide access to Amazon OpenSearch Service resources.	September 08, 2021
AWSQuickSightOpenSearchPolicy – New policy	Amazon Quick added a new policy to allow access to Amazon OpenSearch Service resources from Quick.	September 08, 2021
Amazon Quick started tracking changes	Amazon Quick started tracking changes for its AWS managed policies.	August 2, 2021

Ask questions, explore data, and get insights with chat in Amazon Quick

You can use Amazon Quick chat to ask questions, interact with your Quick resources, and accomplish tasks using multiple agentic workflows using natural language. When you ask a question, Quick chat analyzes the data it has access to to generate comprehensive responses.

Quick chat provides you with the following capabilities.

Topics

- [Welcome message](#)
- [Suggested prompts](#)
- [Expand and collapse modes](#)
- [Agent picker](#)
- [Scoping response generation to specific data](#)
- [Chatting with datasets](#)
- [Chatting with Topics](#)
- [Web search](#)
- [Upload files and chat](#)
- [Actions](#)
- [Flows](#)
- [Memory and response personalization](#)
- [Source citations](#)
- [Copy responses](#)
- [Provide feedback](#)
- [Artifact panel](#)
- [Conversation management](#)
- [Response events](#)
- [Explanations](#)
- [Contextual awareness](#)
- [Managing your preferences in Amazon Quick](#)
- [Amazon Quick chat explanations](#)

Welcome message

When you log in, Amazon Quick chat displays a welcome message.

Suggested prompts

To help you get started on interacting with your resources in chat, Amazon Quick displays suggested prompts you can use to learn more about Amazon Quick features. Use these to get started on your tasks. Once you select a prompt, Amazon Quick guides you step-by-step through the information it needs from you to complete the task.

Note

Suggested prompts are system generated for the system agent; you can't customize them. For custom agents, suggested prompts can be created and customized by the agent builder or owner.

Expand and collapse modes

To help enhance your productivity and streamline your chat interactions, Amazon Quick offers two separate chat window display modes:

- **Collapse mode** – The default chat window display. Use for easy access to chat from within all Amazon Quick contexts.
- **Expand mode** – Access chat in an expanded view. Use this mode to work in the chat window without any distractions.

You can open up the chat interface at any time by selecting the chat bubble icon from the top right navigation menu. You can also exit out of the chat at any time by selecting the cancel icon from the top navigation menu.

Agent picker

Chat agents in Amazon Quick are conversational assistants that help users explore data, analyze information, and take actions. You can pick any agent you've created to use during chat, or use the system agent **My Assistant** for your conversations.

To learn more about agents, see [Working with agents](#).

Scoping response generation to specific data

When you use Amazon Quick, you can define what data you want your chat responses generated from using data filters. Data filters enable you to customize your chat's expertise area. The data available to your agent depends on whether it is bound to specific Amazon Quick knowledge and integrations, or if it's unbounded (not connected to specific Amazon Quick resources).

By default, an agent has knowledge of all your Amazon Quick resources and its underlying large language model (LLM). While using chat, you can select between the following knowledge modes:

- **All data and apps** – Amazon Quick will use all available knowledge sources (including spaces, dashboards, datasets, topics, knowledge bases, action integrations, and LLM knowledge) to generate a response to your prompts.
- **General knowledge** – Amazon Quick will only use LLM knowledge to generate responses to your prompts. It won't use any knowledge from Amazon Quick resources.
- **Specific data and apps** – Chat with a specific or multiple spaces, dashboards, datasets, topics, knowledge bases, or actions. To choose specific resources, select **Choose resource**.

An agent with pre-defined knowledge, on the other hand, only has knowledge of existing Amazon Quick resources already linked to it, and its underlying LLM knowledge. While chatting with an agent with pre-defined knowledge, you get the following knowledge modes:

- **Pre-defined knowledge** – Amazon Quick will use only the Amazon Quick resource knowledge linked to the agent, and LLM knowledge, to generate responses. You can't change this setting.
- **Additional data and apps** – Chat with a specific or multiple spaces, dashboards, datasets, topics, or knowledge bases and actions. To choose specific resources, select **Choose resource**.

Chatting with datasets

You can chat with your datasets in Amazon Quick to ask questions and get insights from your data using natural language. To get started, you need to first create a dataset. For more information, see [Creating datasets](#).

Amazon Quick chat supports the following dataset types:

- SPICE datasets
- Direct Query datasets against Amazon Redshift, Athena, Aurora PostgreSQL, and Amazon S3 Tables

 Note

Direct Query datasets against other data sources are not yet supported for chat.

For SPICE datasets, the size limit of 2 TB and 2 billion rows applies when chatting with the dataset.

Chatting with Topics

You can chat with your Topics in Amazon Quick to ask natural language questions that span multiple datasets. When you select a Topic as your data context in chat, the LLM-powered chat agent uses the defined relationships, dataset enrichment metadata, and custom instructions to generate cross-dataset SQL queries and return unified answers.

To chat with a Topic, use the data filter to select **Specific data and apps** and choose your Topic from the list. Or navigate to the Topic and choose the chat icon.

The chat agent traverses relationships across datasets, identifying which tables contain the relevant columns, constructing SQL with appropriate joins, and returning a unified answer. You can view the generated SQL in the Explanation panel to verify the correct datasets and joins were used.

Amazon Quick chat supports the following Topic types:

- **New Topics (multi-dataset)** – The LLM-powered chat agent generates cross-dataset SQL with runtime joins. It can produce inner joins, left joins, outer joins, unions, and subqueries based on your defined relationships and custom instructions.
- **Legacy Topics** – The legacy ML-based fuzzy search model selects one dataset from the Topic and queries only that single dataset. Cross-dataset queries are not supported.

For more information about creating and configuring Topics, see [Working with Amazon Quick Sight Topics](#).

Web search

Amazon Quick comes with built-in web search capability. Amazon Quick chat automatically searches both your connected data sources and the web, providing comprehensive answers from either or both sources as needed. You can turn web search off for any query by clicking the globe icon in the chat panel. Any content retrieved from the web is appropriately cited in the final response.

Note

Web search is available to users only if it has been enabled by admin.

Upload files and chat

You can upload documents into Amazon Quick chat and ask Amazon Quick chat to summarize or analyze data based on the content of the uploaded documents.

When you start a new conversation, you can upload new files, choose from a saved list of recent documents, or drag and drop files directly into the conversation.

You can upload up to 20 files at one time with the following size limits:

- 10 MB for images (.jpeg, .png)
- 10 MB for spreadsheet files (.csv, .xls, .xlsx)
- 50 MB for other supported formats

Supported file types include:

- Word (.docx)
- Excel (.xls, .xlsx)
- PowerPoint (.ppt, .pptx)
- PDF (.pdf)
- images (.jpeg, .png)
- text files (.csv, .txt, .rtf, .md)
- Outlook (.msg)

- structured data files (.json, .yaml, .xml, .vtt, and common code files)

The total parsed content for all files combined has to be under 665,000 characters.

Note

If your files exceed the file count and character count limits, you can upload them to a [space](#) instead.

Documents uploaded through the chat interface are deleted with the associated conversation after 90 days of inactivity.

Note

Audio and video files are not supported as file uploads.

Actions

Amazon Quick actions boost productivity by allowing you to perform relevant tasks in supported third-party services from your Amazon Quick chat interface, like sending an email using Microsoft Outlook, or creating a PagerDuty incident. You can invoke any actions configured and available to you from your Amazon Quick chat interface. To select an action, use the action application from the data scoping filter and select the third party application you want to perform actions in. Then, choose the action you want to perform from the available actions. The system checks if you have authorization to perform the action before performing it.

To learn more about actions, see [Actions in Amazon Quick](#).

Flows

Use Amazon Quick flows to create shortcuts for everyday, repetitive tasks you perform while chatting in Amazon Quick. For example, if you find yourself using Amazon Quick chat to create and copy edit product launch descriptions, you can use Amazon Quick flows to generate an application that takes your draft content as input, and generates a polished draft—after content, style, and grammar review—as output.

You can run a flow directly from chat within a conversation. Within a single conversation, you can talk to a chat agent as well as run a flow.

To learn more about flows, see [Flows in Amazon Quick](#).

Memory and response personalization

Memory in Amazon Quick chat helps personalize your conversations by remembering your preferences, context, and past interactions. Chat agents can reference this stored information to provide more personalized and contextually relevant responses without you needing to repeat information in every conversation, enabling more relevant and tailored responses over time.

When you interact with chat agents in Amazon Quick, you can provide your preferences including:

- Response style and format
- Frequently accessed Dashboards, Spaces, integrations, and other Amazon Quick assets
- Context about your role and responsibilities
- Common tasks and workflows you perform

This information is securely stored and fetched to enhance your future interactions with chat agents. Memories are specific to you and are not shared with other users. To get the most value from memory in chat, follow the best practices below:

- **Be explicit about your preferences** – Clearly state your preferences, priorities, and context when you want the chat agent to remember them. For example, "I prefer reports in table format" or "I work on the marketing team".
- **Provide relevant context** – Share information about your role, responsibilities, and common tasks to help chat agents understand your needs better.
- **Update your preferences** – As your priorities change, inform the chat agent about updates to ensure the stored information remains current and relevant.
- **Review and refine** – If responses don't align with your expectations, provide feedback and clarify your preferences to improve future interactions.

You maintain control over your memory data. You can review chat memory by going to **Conversation History** and clicking on **View Memories**.

 Note

Chat memory has the following limitations:

- Chat memory isn't supported for accounts configured with [AWS KMS customer managed keys](#).

Source citations

Amazon Quick chat responses provide in-text source citations in a numbered list. To view the source for a response, choose the number at the end of the sentence. The dialog box shows the title, URL, and a snippet from the source that was used to generate the response. Choose the URL to view the source document.

To view the full list of sources, choose **Sources** at the end of the response. You can use the source list to fact-check the response, or for deeper analysis.

Copy responses

You can copy and save the responses for later review and analysis. To copy a response, choose the copy icon at the end of the response.

Provide feedback

You can provide direct feedback about the responses you received from Amazon Quick chat by using the thumbs-up or thumbs-down button. Your feedback is used to help address technical issues in the web experience.

If you select the thumbs-down button, you can choose from the following feedback options:

- **Inaccurate, not factually correct**
- **Incomplete answer**
- **Didn't understand my question**
- **Issue with sources**
- **Too wordy**

- **Offensive or unsafe language**
- **Other**

You can also copy conversation details from the feedback window for further troubleshooting.

Artifact panel

An *artifact* in Amazon Quick is defined as any Amazon Quick generated content produced using Amazon Quick's underlying large language model's (LLM) generative capabilities, for example: contracts, email drafts, social media posts, tables, interactive charts, code blocks, or any Amazon Quick visuals.

Amazon Quick stores any artifacts it generates as responses during chat in an artifact panel. You can access any artifacts created during your chat by selecting **View more** from below the artifact in chat.

You can download artifacts as text, CSV, code files, Word documents, Excel spreadsheets, PowerPoint presentations, PDFs, and infographics. Saving artifacts to a Space is not currently supported.

For more information about creating and downloading documents and visuals, see [Document and visual creation](#).

Conversation management

Amazon Quick stores conversations for up to 90 days, and you can access them in the left navigation pane. You can perform the following tasks to manage your conversations:

- **Start new conversation** – Choose **+** from the chat controls menu to start a new conversation.
- **View conversation history** – Choose the clock icon to view all conversation history for the last 90 days. You can filter conversation history to specific agents.

Note

Searching, saving, renaming, and sharing conversations with other users is not supported.

Response events

For every natural language prompt Amazon Quick chat receives, it generates real-time processing steps while the request is being fulfilled. These steps can be viewed as part of **Response events** during response generation, and are also accessible once Amazon Quick has finished generating a response. Response events vary based on the chat interaction you're having (whether you're chatting with a file you've uploaded or with a Amazon Quick resource) and capture any rewrites to your response before it's finalized. A response event also includes the system generated identifiers for your conversation—you can copy and save the information for troubleshooting..

Explanations

When you chat with dashboards and datasets, each answer includes an Explanation that shows how the model arrived at each numerical claim, including the data sources, assumptions, filters, calculations, and SQL queries that the model used. Instead of manually verifying each answer by finding the original source and re-creating the logic, you can directly see the model's assumptions at the click of a button.

When you open an Explanation, you can see the following components:

- **Found data in** – Displays the dashboards, sheets, or datasets where the insight came from.
- **Filters** – Lists dashboard filter values that were used to arrive at the answer.
- **Assumptions** – Unpacks any large language model (LLM)-derived definitions from either the data directly or from world knowledge.
- **Calculation explained** – Shows any calculations that the model performed to arrive at the answer, presented in both natural language and as a math formula.
- **Generated SQL** (datasets only) – Displays the specific SQL query that produced each numerical claim.

For more information and examples, see [the section called “Chat explanations”](#).

Contextual awareness

Amazon Quick chat is context aware. It registers the Amazon Quick resource you're interacting with and scopes your chat automatically to the resource. This way, when you start chatting, the resource you're interacting with is already included as a knowledge source for your chat.

If you navigate to a specific Amazon Quick resource during an on-going chat, and that resource supports chat, Amazon Quick chat will automatically register your new context and ask you whether you want to chat with the specific resource you have navigated to. If you want to continue chatting in the mode you're in, you can choose to **Ignore** the message. Or you can choose **Yes** to focus your chat on the specific resource you're interacting with. If you do so, Amazon Quick automatically scopes your chat interactions to the resource in context.

Managing your preferences in Amazon Quick

User Preferences in Amazon Quick gives you control over how Quick looks, feels, and works for you. You can customize your chat experience, set a default chat agent, personalize how Quick interacts with you, and manage your memories — all from a single settings modal.

Customizing your chat experience

You can set your default chat panel to open in either expanded or collapsed mode. By default, Quick automatically remembers your last used setting. For example, if you expand the chat panel during a session, Quick updates your default so the panel opens expanded the next time you return. You can also manually set your preferred default at any time from the Preferences modal.

Setting your default chat agent

You can choose a preferred chat agent that greets you each time you open Quick. In addition to selecting a default agent, you can pre-select a default knowledge scope for My Assistant. The available knowledge scope options are:

- **All data and apps** – Quick searches across all your connected data sources and applications.
- **General knowledge** – Quick responds using general knowledge only, without searching your connected data.
- **Specific data and apps** – Quick searches only the data sources and applications you specify.

Your selected agent and knowledge scope are applied automatically each time you start a new conversation.

Personalizing your experience

You can let Quick know what to call you and share your area of focus at work. Quick uses this context to personalize responses and make your experience more relevant to your role and responsibilities.

Managing your memories

You can view and manage your memories directly from Preferences. Memories are insights that Quick learns from your conversations to personalize future interactions. From the Memories section, you can:

- View all saved memories.
- Search for specific memories.
- Delete individual memories.
- Delete all memories at once.


Accessing User Preferences

To access and configure your preferences:

1. Go to the top-right corner of the Amazon Quick homepage.
2. Choose your account name.
3. Select **Preferences** from the dropdown menu.
4. Make your changes across any of the available sections.
5. Choose **Save** to apply your changes, or choose **Cancel** to discard them.

Amazon Quick chat explanations

In Amazon Quick, when you chat with dashboards and datasets, each answer includes an Explanation that shows how the model arrived at each numerical claim, including the data sources, assumptions, filters, calculations, and SQL queries that the model used. Instead of manually verifying each answer by finding the original source and re-creating the logic, you can directly see the model's assumptions at the click of a button.



Chat with your dashboards

When you chat with your dashboard data, open the Explanation to see which dashboards and sheets were selected. You can also see which filters were applied. This helps you validate whether the answer matches your intent.

For example, suppose you open your "Test Drive Conversion" dashboard and ask "what electric car models have an almost perfect satisfaction score but a low conversion rate?" You want to see if some cars test drive well but don't result in a sale. You open the Explanation and check the Assumptions section. Chat defined "electric car model" by using the vehicle model naming terminology. It searched for names that ended with "E" (electric) or "SE" (sport electric). Although that might be correct in most cases, the best field to ensure accuracy is "vehicle_fueltype". You type that directly in the chat: "use the vehicle fuel type for identifying electric cars." You then open the refreshed and correct Explanation.

ⓘ Explanation

×

📄 Data point 1

“Table describing electric vehicle models with satisfaction scores and conversion rates”

^

Found data in



Test Drive Performance and Conversi...

4 filters ^

Vehicle & Engagement Analysis

Modified

Existing

Categorical

status equals: all

customer_role equals: all

country equals: all

vehicle_fueltype equals: all

Assumptions

- *'electric car models'* → identified models ending with ' E', ' SE', or containing 'SE ALL4'
- *'satisfaction score'* → mapped to *avg_satisfaction* column from 'Average Customer Satisfaction by Vehicle Model' visual
- *'conversion rate'* → calculated from *lead_to_offer* column in 'Lead Conversion by Vehicle Model' visual

Calculations explained

Retrieved satisfaction scores from the 'Average Customer Satisfaction by Vehicle Model' visual and conversion data from the 'Lead Conversion by Vehicle Model' visual on the Vehicle & Engagement Analysis sheet. Identified electric models by filtering for model names ending in ' E' or ' SE' or containing 'SE ALL4'. Calculated conversion rate for each model by dividing the count of test drives that converted (*lead_to_offer* = true) by total test drives and multiplying by 100.

▼ Calculation

📄 Data point 2

▼

Explanation



Data point 1



"Table describing all 6 electric car models with satisfaction and conversion metrics"

Found data in

Test Drive Performance and Conversi...
4 filters ^

Vehicle & Engagement Analysis

Modified Existing

Categorical

status equals: all customer_role equals: all

country equals: all vehicle_fueltype equals: ELECTRIC

Assumptions

- 'electric car models' → filtered column *vehicle_fueltype* to *ELECTRIC*
- 'satisfaction' → used column *feedback_rating* from visual 'Average Customer Satisfaction by Vehicle Model'
- 'conversion rate' → calculated from column *lead_to_offer* from visual 'Lead Conversion by Vehicle Model'

Calculations explained

Filtered the Vehicle & Engagement Analysis sheet to show only electric vehicles using the *vehicle_fueltype* filter. Retrieved satisfaction scores from the 'Average Customer Satisfaction by Vehicle Model' visual and conversion data from the 'Lead Conversion by Vehicle Model' visual. Calculated conversion rate for each model by dividing the count of converted test drives by total test drives and multiplying by 100.

Data point 2



"The average satisfaction across electric models is **2.991106** and the average conversion rate is **34.779935%**."

Data point 3



"MINI Aceman E — Satisfaction of **3.279070** (2nd highest) but conversion of only **34.285714%**, just below the electric average. Despite 105 test drives

Explanation components

- Found data in** – Displays the dashboards and corresponding sheets where the insight came from.
- Filters** – Lists dashboard filter values that were used to arrive at the answer.

- **Assumptions** – Unpacks any large language model (LLM)-derived definitions from either the data directly (like referencing agent instructions) or from world knowledge.
- **Calculation explained** – Shows any calculations that the model performed to arrive at the answer, presented in both natural language and as a math formula.

Chat with your datasets

When you chat directly with your datasets, you can see the SQL queries that are generated. Use these queries to verify that the model understood your intent. In the car dealer example, suppose you ask "what's the no-show rate, and which car model struggles with it the most?"

- **Found data in** – Displays the datasets where the insight came from.
- **Assumptions** – Unpacks any LLM-derived definitions from either the data directly (like referencing dataset-level descriptive metadata) or from world knowledge.
- **Calculation explained** – Shows any calculations that the model performed to arrive at the answer, presented in both natural language and as a math formula.
- **Generated SQL** – Displays the specific SQL query that produced each numerical claim.

🔍 Explanation



💬 Data point 1



"Out of **2,000** total test drives, **338** were no-shows — giving an overall no-show rate of **16.90%**."

Found data in

 VBS_test_drives

Assumptions

- '*no-show rate*' → filtered column *status* to 'NO_SHOW' and calculated percentage

Calculations explained

Queried the VBS_test_drives dataset to count total test drives (2,000) and those with status 'NO_SHOW' (338), then calculated the percentage: 338 divided by 2,000 equals 16.90%.

^ Calculation



$$\text{No-Show Rate} = \frac{338}{2,000} \times 100 = 16.90\%$$

^ SQL query



```
SELECT
  COUNT(*) AS total,
  SUM(
    CASE
      WHEN "status" = 'NO_SHOW' THEN 1
      ELSE 0
    END
  ) AS no_shows,
  ROUND(
    100.0 * SUM(
      CASE
        WHEN "status" = 'NO_SHOW' THEN 1
        ELSE 0
      END
    ) / COUNT(*),
    2
  ) AS no_show_rate_pct
FROM
  "VBS_test_drives"
```

Document and visual creation with Amazon Quick

Quick can create professional-quality documents and business visuals from natural language requests. You can create and edit Word documents, Excel spreadsheets, PowerPoint presentations, PDFs, infographics, and other business visuals directly within Quick. Describe what you need, and Quick produces a downloadable file that you can use immediately or refine further.

Document and visual creation is available in the following Quick features:

- **Quick chat** — Create and edit documents and visuals through conversational requests. For more information, see [the section called “Creating documents and visuals in chat”](#).

Topics

- [Supported document and visual types](#)
- [How document and visual creation works](#)
- [Creating documents and visuals in chat](#)
- [Editing existing documents](#)
- [Working with uploaded files](#)
- [Using templates](#)
- [Previewing and downloading](#)
- [Business visual creation](#)
- [Brand theming](#)
- [Best practices](#)
- [Limitations](#)

Supported document and visual types

Quick supports creating the following document and visual types:

Word documents (.docx)

Create professional reports, proposals, contracts, and other text-based documents. Quick generates formatted documents with styles, tables, images, charts, headers, footers, and tables of contents.

Excel spreadsheets (.xlsx)

Create spreadsheets with formulas, charts, conditional formatting, and multiple worksheets. Quick supports dashboards, financial models, data analysis workbooks, trackers, and data exports.

PowerPoint presentations (.pptx)

Create slide decks with professional themes, layouts, charts, speaker notes, and visual elements. Quick can create presentations from scratch or from an existing template.

PDF documents (.pdf)

Create formatted PDF documents with styled pages, tables, headers, footers, and page numbers.

Business visuals (.png)

Create visual infographics, posters, diagrams, and charts from natural language descriptions.

How document and visual creation works

When you request a document or visual, Quick uses an AI-powered process to interpret your request and produce the file. The process works as follows:

1. You describe the document you want to create using natural language. For example, *"Create a quarterly sales report in Excel with charts showing revenue by region."*
2. Quick determines the appropriate document type and generates the file using specialized document creation tools.
3. As the document is being built, you see real-time progress updates describing each step of the generation process.
4. When the document creation is complete, you can preview it in a dedicated panel to review the content and format.
5. If everything looks good in the preview, you can download the document. If you need to make edits and refine the content further, you can add comments to specific sections of the document and regenerate the document with the feedback included. For broader changes that apply across the document, you can also use chat to give feedback and suggested changes.

You can also upload files to use as input for document generation. For example, you can upload a CSV file and ask Quick to create an Excel dashboard from the data, or upload a Word document and ask for specific edits.

Creating documents and visuals in chat

You can create documents and business visuals directly in Quick chat by describing what you need.

Create a new document or visual

1. Open Quick chat.
2. (Optional) Upload any input files that you want to use as source data. You can upload spreadsheets (.csv, .xls, .xlsx), documents (.docx), presentations (.pptx), PDFs, and JSON files.
3. Describe the document you want to create. Be specific about the content, structure, formatting, and any analytical requirements. See the following section for examples.
4. Quick generates the document and displays progress updates as it works. When the document is ready, it appears in a preview panel.
5. Choose **Download** to save the file to your device.

Example prompts

The following examples demonstrate the range of documents and visuals you can create, from simple documents to complex multi-sheet financial models.

Financial modeling (Excel)

"Our company is evaluating whether to build a new distribution center or lease additional warehouse space. Create an Excel model comparing the total cost of ownership for building versus leasing over 10 years. Include an assumptions sheet with construction costs, lease rates, annual escalation percentages, maintenance costs, and discount rate that I can adjust. Create a 10-year cash flow projection for each scenario with NPV and IRR calculations. Add a break-even analysis sheet and a conclusion page comparing the two options with charts."

This produces a multi-sheet workbook with an editable assumptions page, year-by-year cash flow projections for each scenario, Excel formulas that update when you change inputs, and a conclusion sheet with comparison charts and financial metrics.

Data-driven presentation (PowerPoint)

"Using the attached quarterly sales CSV, create a board-ready PowerPoint presentation. Include an executive summary slide with key metrics, regional breakdowns with bar charts, a year-over-year trend analysis with line charts, and a final slide with strategic recommendations. Use a dark professional theme with our company colors (navy and gold)."

Upload your data file alongside the prompt. Quick extracts the data, builds the charts, and applies the requested theme across all slides.

Contract document (Word)

"Create a professional services agreement in Word. Include sections for scope of work, payment terms with a milestone schedule table, intellectual property rights, confidentiality, termination clauses, and limitation of liability. Add a signature block at the end. Use formal legal formatting with numbered sections and sub-sections."

This produces a structured legal document with proper heading hierarchy, numbered clauses, formatted tables, and signature blocks.

Operational dashboard (Excel)

"Create a project management dashboard in Excel. Include a summary sheet with KPI cards for budget utilization, schedule variance, and resource allocation. Add a task tracker sheet with status columns that use conditional formatting (green for on track, yellow for at risk, red for blocked). Add a Gantt-style timeline sheet and a risk register with severity scoring formulas."

This produces a multi-sheet workbook with conditional formatting, formulas that calculate status indicators, and charts that visualize project health.

Research report (Word with data)

"Using the attached PDF research paper and the CSV dataset, create a Word report that summarizes the key findings, includes a methodology section, presents the data in formatted tables with alternating row shading, and adds charts showing the main trends. Include a table of contents and an executive summary at the top."

Upload both files with your prompt. Quick reads the PDF content and CSV data, then produces a formatted report that synthesizes both sources with embedded charts and tables.

Visual infographic

"Create an infographic showing the global impact of renewable energy adoption. Include statistics on solar and wind capacity growth over the last decade, a world map highlighting top adopting

countries, and a comparison of energy costs between renewable and fossil fuel sources. Use a clean, modern design with a green and white color palette."

This produces a visually designed infographic with data visualizations, geographic elements, and styled typography. Because no source data is uploaded with this prompt, the statistics shown are illustrative. You can refine the values by providing a data file or describing updates in chat.

Editing existing documents

You can edit documents that Quick has previously created, or upload your own files for editing. Quick preserves the existing formatting and structure while applying your requested changes. For most document types, you can upload your own files for editing. See individual document types below for details.

Edit an existing document

1. In Quick chat, upload the document you want to edit, or continue a conversation where a document was previously generated.
2. Describe the changes you want to make. For example:
 - *"Change all the heading colors to dark blue."*
 - *"Add a new column called Status to the first table."*
 - *"Update the title slide to say Q1 2026 Results."*
3. Quick applies the edits and produces an updated version of the document.

The following editing capabilities are supported for each document type.

Word (.docx)

- *Text replacement* — Replace words or phrases throughout the document. Example: *"Change all instances of 2025 to 2026."*
- *Formatting changes* — Update fonts, colors, sizes, and styles. Example: *"Make all Heading 1 text dark blue and 16pt."*
- *Section management* — Add, remove, or reorder sections. Example: *"Move the Conclusions section to before the Appendix."*

- *Table modifications* — Add rows, columns, or update cell values. Example: *"Add a Status column to the project tracking table."*
- *Headers and footers* — Add or update page headers, footers, and page numbers. Example: *"Add a footer with the company name and page numbers."*

Excel (.xlsx)

- *Cell updates* — Modify values, formulas, and number formats. Example: *"Update the Q4 revenue figure in cell B5 to \$3.5M."*
- *Chart additions* — Add or modify charts and visualizations. Example: *"Add a bar chart comparing revenue across all quarters."*
- *Worksheet management* — Add new sheets or restructure existing ones. Example: *"Add a Summary sheet with totals from each regional sheet."*
- *Formatting* — Update styles, conditional formatting, and layout. Example: *"Apply alternating row colors and bold the header row."*

PowerPoint (.pptx)

- *Content updates* — Change text, titles, and bullet points. Example: *"Update the title slide to say Q1 2026 Results."*
- *Layout changes* — Modify slide layouts and visual elements. Example: *"Change slide 3 to a two-column layout with the chart on the right."*
- *Data modifications* — Update chart data and table content. Example: *"Update the sales chart with the new figures from the uploaded spreadsheet."*
- *Slide management* — Add, remove, or reorder slides. Example: *"Add a new slide after the introduction with a team overview."*

Infographics (.png)

Infographic editing applies to visuals that were created in the current conversation. You can request changes to a previously generated infographic without uploading a separate file.

- *Text changes* — Update titles, labels, and descriptions. Example: *"Change the title to Global Warming Trends."*
- *Visual updates* — Modify colors, layout, and styling. Example: *"Make the background blue and increase the font size of the statistics."*

Note

PDF documents cannot be edited in place. To modify a PDF, Quick extracts the content and creates a new PDF with your requested changes.

Working with uploaded files

You can upload files to use as input when creating documents and visuals. Quick reads the content of your uploaded files and uses it to produce the requested output.

Common use cases include:

- Uploading a CSV or JSON file and creating an Excel dashboard or report from the data
- Uploading a PDF and converting its content into a Word document or presentation
- Uploading multiple source files and creating a consolidated summary document
- Uploading a PowerPoint template and creating a new presentation that matches its visual style

Supported input file types:

- Spreadsheets: .csv, .xls, .xlsx
- Documents: .docx, .pdf
- Presentations: .pptx
- Data files: .json
- Brand assets: .svg (for logos used with brand theming)

Using templates

You can upload an existing file as a template so that Quick creates new content that matches the original's visual design, branding, and formatting. Template support varies by document type, with PowerPoint offering the most comprehensive template cloning workflow. This is useful when your organization has standardized document styles that you want to maintain across new documents.

PowerPoint templates

PowerPoint has the most comprehensive template support. When you upload a .pptx file and ask Quick to create a new presentation using it as a template, Quick does the following:

1. Inspects the template to discover slide layouts, shape names, theme fonts, and theme colors.
2. Classifies the template slides by type, such as title or cover slides, section dividers, and content slides.
3. Clones the template slides that best match each section of your new content. Cloning preserves backgrounds, brand colors, decorative elements, icons, and images from the original template.
4. Fills in your new content by replacing text in the cloned slides while keeping the visual design intact.

The result is a new presentation with your content that looks like it was built by hand using the original template.

To use a PowerPoint template:

1. Upload your .pptx template file in Quick chat.
2. Describe the new presentation you want to create. Mention that you want to use the uploaded file as a template. For example:
 - *"Using this template, create a 12-slide presentation about our 2026 product roadmap."*
 - *"Create a quarterly business review deck using this template and the attached CSV data."*
3. Quick generates the presentation using the template's visual design and your new content.

PowerPoint templates work best when the template contains a variety of slide layouts, such as title slides, content slides with text placeholders, and section dividers. Quick selects the most appropriate layout for each slide in the new presentation based on the content.

Excel templates

When you upload an existing .xlsx file as a template, Quick clones the workbook and preserves its formatting, formulas, and structure. It then populates the cloned workbook with your new data while keeping the original styling intact.

This is useful for:

- Recurring reports that use the same layout each period, such as monthly financial statements or quarterly dashboards
- Standardized workbooks with pre-built formulas, conditional formatting, and data validation rules that you want to reuse with new data
- Branded spreadsheets with specific header styles, color schemes, and chart formatting

To use an Excel template, upload your .xlsx file and describe what data to populate. For example: *"Using this template, fill in the revenue sheet with the data from the attached CSV and update the dashboard charts."*

Word templates

Word documents do not have a dedicated template cloning workflow. However, you can achieve a similar result by uploading an existing .docx file and asking Quick to edit it with new content. When editing, Quick preserves the original formatting, styles, headers, footers, and structure while applying your requested changes.

For example, if your organization has a branded report template with specific fonts, colors, and section layouts, you can upload it and ask: *"Replace the placeholder content in this report template with the following project details..."*

For new Word documents created from scratch, Quick applies professional formatting with configurable styles. You can specify fonts, colors, and layout preferences in your prompt to match your organization's standards.

Previewing and downloading

When Quick creates a document or visual, it appears in a preview panel. You can access the preview panel by choosing **View more** from below the item in chat.

From the preview panel, you can:

- Preview the created document or visual
- Download it to your device
- Continue editing by sending follow-up messages in chat

The following download formats are supported:

Type	Download format
Word documents	.docx
Excel spreadsheets	.xlsx
PowerPoint presentations	.pptx
PDF documents	.pdf
Infographics and business visuals	.png

Business visual creation

Quick can create business visuals like infographics, charts, and data visualizations. You can turn insights into shareable images for presentations and reports, or generate professional graphics for business communication.

Business visuals are available in two ways:

Standalone infographics and visuals

You can create standalone business visuals by describing what you need in Quick chat. Quick produces a professionally designed image that you can download and share. These visuals are ideal for:

- Data-driven infographics that summarize key metrics and trends
- Comparison charts and diagrams for stakeholder communication
- Visual summaries for executive briefings and team updates
- Process diagrams and workflow visualizations

Example prompts:

- *"Create an infographic comparing our Q1 and Q2 performance metrics with charts showing revenue growth, customer acquisition, and retention rates."*

- *"Create a visual summary of our three-year strategic roadmap with milestones and key deliverables for each phase."*
- *"Using the attached CSV, create a data visualization showing regional sales distribution with a map and supporting bar charts."*

You can also edit previously created visuals by describing the changes you want. For example: *"Change the color scheme to blue and white and update the title to say H1 2026 Results."*

Visuals embedded in documents

When creating Word documents, PowerPoint presentations, or PDFs, you can request that Quick include generated visuals within the document. These embedded visuals include:

- Background images and hero visuals for presentation title slides and section dividers
- Illustrations and concept visuals for reports and proposals
- Custom icons and decorative elements with transparent backgrounds
- Data charts and visualizations created from your uploaded data

To include visuals in a document, mention them in your prompt. For example:

- *"Create a PowerPoint presentation about our product launch with generated hero images on the title and section divider slides."*
- *"Create a Word proposal with an illustrated cover page and concept visuals for each solution section."*

Note

Each visual takes a few seconds to generate. For best results, limit visual generation requests to a few images per document.

Brand theming

You can apply your organization's brand identity to created documents and visuals by providing theme assets. When brand theme files are available, Quick automatically applies your brand colors and logos.

Brand theming supports the following asset types:

- Theme configuration files that define brand colors
- Logo files in SVG format

To apply brand theming, upload your theme configuration file (JSON format) and SVG logo assets in Quick chat alongside your document request. The theme configuration file defines your brand colors as hex values. Quick detects the theme files and applies your brand colors and logos to the generated output.

 Note

Brand theming for business visuals and infographics supports brand colors and logos. Typography and additional brand image assets are not included in theme configuration but can be provided as file uploads in chat.

Brand theming is available for all document and visual types, including Word documents, Excel spreadsheets, PowerPoint presentations, PDFs, and infographics.

Best practices

Follow these best practices to get the best results:

- **Be specific in your requests.** Include details about the content, structure, formatting, and any specific requirements. For example, instead of "Create a report," try "Create a Word report with an executive summary, three data sections with tables, and a conclusion."
- **Provide source data when possible.** Upload CSV, Excel, or JSON files so that Quick can use real data instead of generating placeholder content.
- **Iterate on your documents.** Start with a first draft and then request specific edits in follow-up messages. This is often more effective than trying to specify everything in a single request.
- **Specify the document type.** Mention the file format you want (Word, Excel, PowerPoint, PDF) in your request to ensure Quick generates the correct document type.
- **Use templates for consistent styling.** When creating presentations, upload an existing .pptx template to maintain consistent visual styling across your organization.

Limitations

The following limitations apply to document and visual creation in Quick:

- When you provide source data, Quick uses it exactly as given and does not fabricate additional data points. When no source data is provided, Quick can generate illustrative values such as common assumptions or sample inputs for models and templates. For editable document types (Word, Excel, PowerPoint), these values are clearly presented as defaults that you can modify directly in the downloaded file. For infographics (PNG), you can request changes by describing the updates you want in chat.
- PDF documents cannot be edited in place. To modify a PDF, Quick extracts the content and recreates the document.
- Complex formatting or highly customized layouts may require multiple iterations to achieve the desired result.
- Saving generated documents to a Space is not currently supported.

Create, customize, and deploy AI-powered chat agents in Amazon Quick

Chat agents in Amazon Quick help users explore data, analyze information, and take actions. Users can interact with chat agents using the Quick chat interface. Chat agents provide assistance through open-ended conversations supported by specific goals, sources of knowledge, and any connected tools. Chat agents can evolve from simple question-answering interfaces to more advanced functions that orchestrate complex workflows.

You can use chat agents to:

- Generate content and provide answers through natural language conversations
- Analyze and summarize information from connected spaces, dashboards, topics, datasets, and uploaded files
- Invoke actions to perform predefined steps for consistent, repeatable outcomes

Note

To learn more about chatting using chat agents, see [Using Amazon Quick chat](#).

Topics

- [Chat agent types](#)
- [Amazon Quick user interaction with chat agents](#)
- [Amazon Quick user permissions for chat agents](#)
- [Custom permissions for chat agents](#)
- [System chat agent](#)
- [Custom chat agents](#)
- [Chat agent context sources and best practices](#)
- [Use a chat agent](#)

Chat agent types

Amazon Quick supports two types of chat agents:

- **System chat agent** – This chat agent ("My assistant") is automatically available to all users by default. The system chat agent serves as a base planner with no inherent data or actions of its own—it dynamically accesses resources available to each user during chat time, allowing it to be tailored to individual user permissions and available content. Admin users control system chat agent settings by assigning specific users as owners so they can customize its persona and other settings.

The system chat agent is enabled with all chat capabilities, including file upload functionality, LLM knowledge access, toxicity and other guardrails, web search. It also includes chat data scoping mechanisms that provide access to spaces, topics, dashboards, knowledge bases, and actions based on user permissions.

- **Custom chat agents** – These chat agents can be created and customized to specific use cases by users with chat agent creation capabilities in Amazon Quick, and shared with anyone. Admins can choose to restrict specific users and groups from creating custom chat agents, while still allowing users to use chat agents via chat. Chat agents return responses scoped to content that their invoking users have permissions to.

Custom chat agents can interact with Amazon Quick resources in the following ways:

- **Preconfigured with resources** – These chat agents use only the configured resources (for example, spaces as knowledge sources, connectors as tools) when looking for answers or orchestrating actions as their default behavior. While chatting, users can attach additional resources as per their need or invoke actions directly. For example, chat agents configured with only spaces containing files can't take actions by default unless users explicitly invoke them. Similarly, chat agents configured with only actions rely on LLM knowledge unless users attach a space or dashboard for enterprise-specific answers. However, users may attach additional spaces or directly invoke other actions they may have access to within the chat interface.
- **Not opinionated at build time** – These chat agents are not initially configured with resources when they are built (spaces, or connectors, or both). Chat determines the resource boundary for the chat agent. For example, if no space is configured for the chat agent, chat will default to all spaces or LLM knowledge until the user makes a change. If no connectors are configured for the chat agent, during chat, all actions a user has access to will be available to the agent until user makes a resource selection that doesn't include actions.

User resource selection during chat affects chat agent behavior. If a user selects a space, the chat agent will only answer with data within that space and take actions available within that space (if it is not opinionated at build time). If a user selects a dashboard, topic, or knowledge base, the chat agent will only answer from that source and not take any actions since the focus is changed to a specific data source. Users are expected to select all relevant resources within chat for comprehensive assistance.

 Note

The system chat agent is an unopinionated chat agent by design. If you want chat agents to access all actions irrespective of data focus, configure chat agents with connectors.

Amazon Quick user interaction with chat agents

The following table shows what you can do with chat agents in the admin console versus as a Amazon Quick user. For more information on which roles can access these features, refer to the [Amazon Quick pricing](#) document.

Access Level	Capabilities
Admin Console	- Assign owners for the system chat agent and custom chat chat agents using Manage assets - Control whether users can create chat agents using Custom permissions - Configure instance-wide blocked words and phrases for all chat agents
Amazon Quick user	- Create and maintain custom chat agents (with appropriate permissions) - Configure chat agent personality and response styles and provide reference documents to inform its behavior - Link chat agents with spaces (with dashboards, datasets, topics) as their knowledge source to look for answers - Attach connectors to use as tools

Access Level	Capabilities
	• Share chat agents with specific users and teams • Interact with chat agents through conversations • Analyze data with chat agent assistance • Receive permissions-aware responses from chat agents

Amazon Quick user permissions for chat agents

What you can do with a chat agent also depends on the permissions you're assigned for it. There are two permission types that users can be assigned:

- **Owner** – Owners can edit, share, and delete the chat agent.
- **Viewer** – Viewers can view and use the chat agent.

 Note

If you don't have access to a linked resource as either a viewer or owner, resources added to the chat agent by another owner appear as "Resource unavailable". You can delete these resources as an owner, but you cannot list or add them because resource-level sharing is required.

Quick admins must give users the permission to create chat agents. For information on which roles can create chat agents, refer to the [Amazon Quick pricing](#) documentation. For information on how to provide access to these features, see [Custom permissions](#) in the Quick Admin Guide.

The following table outlines how user permissions determine what you can do with a Amazon Quick chat agent:

Permissions type	Permissions
Owners	• Can access and configure agent behavior (agent persona settings, reference documents) and resources (space and connectors) • Can customize chat agent details like title, description, and suggested prompts to improve usability

Permissions type	Permissions
	• Can share chat agents with users and groups • Can delete the chat agent
Viewers	• Can't customize the chat agent's details • Can receive responses based on the permissions they have to resources

Custom permissions for chat agents

By default, Amazon Quick enables all new features available for the Amazon Quick account, so that users can access them immediately based on their subscription. You can use [custom permissions](#) to restrict specific features. As an admin, when you create a custom permissions profile, you can create two types of restrictions for chat agents:

- You can completely disable all chat agent functionality for users, including chatting with the default agent, chatting with custom agents, and creating new agents. This can be done by restricting the **Chat agent** capability.
- You can also specifically restrict the ability to create agents without impacting chat with agents. This can be done by creating a custom permission profile and restricting only the feature **Create chat agents**.

Note

If you want your users to be able to chat using the system agent but don't want them to create chat agents, restrict their chat agent creation abilities only.

Agent permissions in Amazon Quick can be configured in various combinations to meet your organization's needs. To help you understand how to configure permissions, the following section provides a use-case driven approach that groups common scenarios:

Use case scenario	What users can do
Agent with all capabilities	• Chat with all default and custom chat agents

Use case scenario	What users can do
(No restrictions on chat agents, spaces, knowledge bases, actions, or flows)	• Create and customize new chat agents • Connect chat agents to spaces, knowledge bases, and actions • Trigger flows from the chat window from existing conversation with a chat agent • Share chat agents with other users • Use all chat features and integrations
Chat-only access to agents (Agent capabilities enabled, but chat agent creation restricted)	• Chat with existing default and shared custom chat agents • Can't create new custom chat agents • Can't edit or update custom chat agents • Full access to chat interface and chat agent library
Limited access to other capabilities (Chat agents enabled, but specific capabilities restricted)	• Chat with system and custom agents but with reduced functionality • Can't attach knowledge bases, spaces, or actions to an agent's configuration, or select these resources during chat • If spaces restricted: Agents configured with spaces will fall back to LLM knowledge only. Users can't view or select spaces in the chat resource selector • If knowledge bases restricted: Users can't select knowledge bases to use in chat. Any knowledge bases present in spaces either attached to the agent or selected by user during chat won't be used to generate responses • If actions restricted: Users can't attach actions to agents or invoke actions from chat. Any actions preconfigured with agents won't be used to execute tasks • If flows restricted: Chat won't show flows anymore, so users can't invoke flows while chatting with an agent

Use case scenario	What users can do
No chat agent access (Agent capabilities completely restricted)	• Can't view or access any chat agents • Agent library and navigation are hidden • Can still access and create other Amazon Quick resources, such as creating spaces for file sharing with teams or flows for structured interactions (as long as those capabilities are not also restricted)

System chat agent

The Amazon Quick system chat agent is automatically created when you sign up for Quick. It serves as the primary interface for all users to interact with their data and perform various tasks within the Quick environment. It comes equipped with default guardrails.

Ownership of the system chat agent is managed by the Quick admin. Admins can disable chatting with chat agents including the system chat agent using [custom permissions](#). If enabled, users with non-admin roles (authors and readers) can view, list, invoke, and use the Quick chat chat agent. Its responses are filtered based on user permissions to assets.

The Quick system chat agent is set to use all chat features out of the box with minimal customization required. While these features are also available to custom chat agents, they must be configured during chat agent creation. The system chat agent includes:

- Default persona as the Quick assistant, with professional tone and response style
- File upload in chat capability enabled
- Large language model (LLM) knowledge chat enabled
- Uses default guardrails configured for chat in admin console
- Access to all spaces, topics, dashboards, knowledge bases, and actions based on user permissions
- Web search capabilities

To learn how to customize a system chat agent as an Admin owner, see [Managing chat agent customization](#). To learn how to control access to the system chat chat agent, see [Manage assets](#).

Custom chat agents

Custom chat agents enable Amazon Quick users to create tailored conversational interfaces for specific business needs. Unlike the system chat agent, which is available to all users with chat permissions by default, custom chat agents can be selectively shared and configured by authorized users.

Users with chat agent creation capabilities can customize the chat agent's personality, response style, and capabilities through a natural language interface, or through the agent builder configuration flow. These users can also configure these chat agents with specific knowledge sources (like spaces) and workflows (like actions) while maintaining control over sharing and access permissions for each of those resources separately. Or, chat agents can be left unlinked for general purpose use which allows end users to use the agent with all or some of their resources.

Quick admins must give users the permission to create chat agents. For information on which roles can create chat agents, refer to the [Amazon Quick pricing](#) documentation. For information on how to provide or restrict access to these features, see [Custom permissions](#) in the Quick Admin Guide.

Topics

- [Create and preview custom chat agents](#)
- [Share custom chat agents](#)
- [Embed custom chat agents](#)
- [Update custom chat agents](#)
- [Manage access to a custom chat agent](#)
- [View chat agents](#)
- [View chat agent details](#)
- [Duplicate custom chat agents](#)
- [Delete chat agents](#)

Create and preview custom chat agents

You can create custom chat agents in Amazon Quick from the **Chat agents** menu on the Amazon Quick home page. You can preview and test your chat agent as you build it. When you create a custom chat agent, you are assigned as its owner by default. The following procedure outlines how to create a custom chat agent.

Using natural language

To create a custom chat agent using a natural language prompt

1. Log in to the Amazon Quick console.
2. From the left navigation menu, select **Chat agents**, and then select **Create chat agent**.
3. In **New chat agent**, do the following:
 - a. In the text input box, enter a natural language description for what kind of chat agent you want to create. This is the prompt Amazon Quick will use you create your chat agent. You can use the suggested sample instructions on the page to create your chat agent as well.
 - b. Then, select **Generate**. Selecting **Generate** takes the instruction or goal provided and expands it into chat agent instructions and configurations. It also scans the user's available resources (spaces and action connectors) and finds the most relevant matches based on the chat agent intent provided. The expanded builder UI is opened with these configurations for the user to review, customize, test, and launch.
4. On the **Configure chat agent** page, make sure your chat agent has the right settings and customize your chat agent further. Click on **Update preview** to ensure all changes are saved before you try out the chat agent in preview.
5. When you're ready, select **Launch chat agent** to publish your custom chat agent to the chat agent library and use it in chat.

Note

- Until you select **Launch chat agent**, your chat agent is not available in the chat agent library. If you exit the creation process without launching, the preview version will be deleted and the chat agent won't be saved.
- When you launch a chat agent, it remains private by default. The chat agent only becomes available in other users' libraries after you share it. Once shared, subsequent edits and launches will publish changes to the same chat agent for all users who have access.

Using builder view directly

To create a custom chat agent using builder mode

1. Log in to the Amazon Quick console.
2. From the left navigation menu, select **Chat agents**, and then select **Create chat agent**.
3. In **Agent Creator**, select **Skip**. This opens up the chat agent configuration page from which you can build and customize your chat agent.
4. Add a name for your custom chat agent. This is the name that your chat agent will be identified by.
5. (Optional) Add a description for your custom chat agent that helps users understand the purpose of the chat agent.
6. (Optional) Select an icon for your chat agent.
7. In **Configure chat agent**, customize the following sections:
 - a. In **AGENT PERSONA**, configure your chat agent's personality, identity, tone, and response style:
 - For **Agent identity** – Define the identity of your chat agent. For example, you can give it instructions on what its name and personality are. Amazon Quick will use defaults if left blank.
 - For **Persona instructions** – Add instructions on how your chat agent interacts with users during chat. For example, you can define what your chat agent's primary tasks are. Amazon Quick uses these to customize your chat agent's persona.
 - b. In **Communication style**, for **Pick a response style preset** – Choose a response style. You can choose between the following response style presets or add custom instructions:
 - **Executive** – Optimized for high-level business communication and strategic insights.
 - **Technical** – Optimized for detailed technical explanations and technical problem-solving.
 - **Creative** – Standard configuration for general-purpose interactions.

Each preset has the following settings:

- For **Tone** – Add a natural language prompt to define your agent's tone. Amazon Quick will use it to customize your chat agent's persona.
- For **Response format** – Add a natural language prompt to define your chat agent's response style. Amazon Quick will use it to customize your chat agent's response style. Define the format of your chat agent's responses, for example: "Use bullet points for lists longer than 3 items".
- For **Length** – Specify the length of your chat agent's responses. Define the length of your chat agent's responses, for example: "Keep answers under 100 words".

When you choose a preset to use, Amazon Quick auto-populates the **Tone**, **Response format**, and **Length** instructions for your chat agent's responses based on the style chosen. You can further customize the existing prompts using natural language, or create new ones.

- c. In **Reference documents**, upload files that remain active in your chat agent's memory to guide all interactions. For more information about how reference documents work with other context types, see [Types of agent context](#).
 - Select **Upload files** or drag and drop your documents to attach files that will guide your chat agent's responses.
 - You can attach documents of .pdf, .txt, .html, .md, .csv, .doc or .docx format. Upto 100,000 characters of text will be extracted and accepted from documents uploaded.

8. (Optional) In **Knowledge sources**, choose between the following options:

- Continue without linking knowledge sources
 - Your chat agent will generate responses from large language model (LLM) knowledge and all Amazon Quick resources the interacting user has access to. You can choose to select a specific knowledge source to focus on during chat. Choose this if you want to create a general purpose chat agent.
- To link specific existing spaces
 - Select **Link**.
 - From the **Link spaces** modal, select the spaces you want to link to your chat agent, and then select **Link**.

The system displays a success message to denote successful linking.

If you link a chat agent to a knowledge source, your chat agent will generate responses from large language model (LLM) knowledge and data from linked resources only. You can add up to 1 more knowledge source to this chat agent as temporary context during chat.

- To create and link a new space
 - Select **Create**.
 - From the **Create space** window that opens, select the assets you want to add to your space and then select **Create**.

If you link a chat agent to a knowledge source, your chat agent will generate responses from large language model (LLM) knowledge and data from linked resources only. You can add up to 1 more knowledge source to this chat agent as temporary context during chat.

 Note

For more information on creating and using spaces, see [Working with spaces in Amazon Quick](#).

The system displays a success message to denote successful space creation. Return to the chat agent creation window.

- In the chat agent creation window, from **Knowledge sources** select **Link**.
- From the **Link spaces** modal, select the space you just created, and then select **Link**.

The system displays a success message to denote successful linking.

9. (Optional) In **Actions**, for **Actions** – Choose between the following options:

- To link specific existing actions
 - Select **Link**.
 - From the **Link action connectors** modal, select the action connectors you want to link to your chat agent, and then select **Next**.
 - In **Actions** select the actions you want to add and then select **Link**.

The system displays a success message to denote successful linking.

- To create a connector and link new actions
 - Select **Create**.
 - From the **Actions** home page, select **New action**.
 - From the **New action** window, for **Sources** select the connector you want to add and then select **Next**.
 - In **Actions** review the actions available and then select **Next**.
 - In **Connection details**, enter the connection details needed and select **Add**.

 Note

For more information on creating and using connectors, see [Actions in Amazon Quick](#).

The system displays a success message to denote successful action addition. Return to the chat agent creation window.

- In the chat agent creation window, from **Actions** select **Link**.
- From the **Link action connectors** modal, select the action connector you just created, and then select **Link**.

The system displays a success message to denote successful linking.

10. In **Customization**, do the following:

- a. For **Welcome message** – Add a welcome message for your chat agent to display to your end user.
- b. For **Suggested prompts** – Add example prompts as conversation starters to inform your end users of the chat agents capabilities.

11. Select **Launch chat agent** to create your custom chat agent.

 Note

- Until you select **Launch chat agent**, your chat agent is not available in the chat agent library. If you exit the creation process without launching, the preview version will be deleted and the chat agent won't be saved.

- When you launch a chat agent, it remains private by default. The chat agent only becomes available in other users' libraries after you share it. Once shared, subsequent edits and launches will publish changes to the same chat agent for all users who have access.

Share custom chat agents

When you create a chat agent, you are assigned as its owner by default. You can choose to share a chat agent you own with other Amazon Quick users using the following permissions:

- *Owner* permissions – User can edit, share, use, and delete the chat agent.
- *Viewer* permissions – User can view and use the chat agent.

You can also define access permissions for a chat agent globally, granting all users access to a chat agent and then defining more granular access permissions for specific users and groups.

The following procedure shows you how to share and assign permissions to chat agents.

To share custom chat agents

1. Log in to the Amazon Quick console.
2. From the left navigation menu, select **Chat agents**
3. Then, from the **Actions** column for the chat agent you want to share, select the menu icon, and then select **Share**.
4. In the **Share chat agent** modal that opens, enter the user or group name you want to share the chat agent with. When the user or group name appears, select it. Repeat the action for all groups or users you want to share the chat agent with. The modal displays the users and groups you're adding.
5. Assign permissions and access information for each user or group you're sharing your chat agent with using the dropdown next to the user or group name. You can assign one of two roles:
 - **Owner** – User can edit, share, and delete the chat agent.
 - **Viewer** – User can view and use the chat agent.
6. Select **Share**.

Embed custom chat agents

You can embed a custom chat agent in an external website or application so that users can interact with it outside of the Amazon Quick console. To embed a chat agent, you must first allow the hosting domain in the Amazon Quick admin settings, and then generate the embed code from the chat agent's share options.

Embedding a chat agent involves two steps:

1. **Allow the hosting domain** – A Amazon Quick admin must add the domain (and optionally its subdomains) where the embedded chat agent will be hosted to the domain allow list.
2. **Generate the embed code** – The chat agent owner or an authorized user generates the embed code from the chat agent's share options.

Allow a domain for embedded chat agents

Before you can embed a chat agent on an external website, a Amazon Quick admin must add the hosting domain to the allow list. You must be a Amazon Quick admin to complete this procedure. This ensures that only approved domains can host embedded chat agents.

To allow a domain for embedding

1. Log in to the Amazon Quick console.
2. Choose the profile icon.
3. Choose **Manage account**.
4. In the navigation pane under **Security**, choose **Manage domains**.
5. In the **Domain** field, enter the domain name where you want to host the embedded chat agent (for example, `https://example.com`).
6. (Optional) Select **Include subdomains** to allow embedded chat agents on all subdomains of the specified domain.
7. Choose **Add**.

The domain now appears in the list of allowed domains. You can edit or delete existing domains by choosing the edit or delete icons next to each domain in the list.

Generate the embed code for a chat agent

After you allow the hosting domain, you can generate the embed code for a chat agent.

To generate the embed code for a chat agent

1. Log in to the Amazon Quick console.
2. From the navigation pane, choose **Chat agents**.
3. From the **Actions** column for the chat agent you want to embed, choose the menu icon, and then choose **Edit**.
4. In the chat agent editor, choose **Share**.
5. In the share options, choose **Share via embed**.
6. Copy the provided embed code and paste it into the HTML of your website or application where you want the chat agent to appear.

The embedded chat agent is now available on your website. Users who visit the page can interact with the chat agent directly.

Update custom chat agents

You can update custom chat agents in Amazon Quick from the **Agents** menu on the Amazon Quick home page. The following procedure outlines how to do so.

Note

You can preview how your chat agent customizations updates perform, as you are configuring the chat agent, in the preview chat interface. Make sure to select **Update preview** before testing a change. Once all changes are tested in preview, select **Launch** to publish the changes to the chat agent currently in use.

To update a custom chat agent

1. Log in to the Amazon Quick console.
2. From the left navigation menu, select **Chat agents**

3. Then, from the **Actions** column for the chat agent you want to share, select the menu icon, and then select **Edit**. This creates a preview version of your chat agent to update and test before publishing.
4. In the **Edit agent**, update your chat agent settings, and then select **Launch**.

Manage access to a custom chat agent

If you're the owner of a chat agent, you can choose to change the user access permissions for a chat agent. You can also remove user and group access to a chat chat agent. The following procedure shows how.

To manage access to custom chat agents

1. Log in to the Amazon Quick console.
2. From the left navigation menu, select **Chat agents**
3. Then, from the **Actions** column for the chat agent you want to share, select the menu icon, and then select **Share**.
4. In the **Share chat agent** modal that opens, select the user and group you want to change permissions for. When the user or group name appears, select it and change the permissions. Repeat the action for all groups or users you want to change permissions for. The modal displays the users and group permissions you're changing. You can take the following three actions:
 - Change to **Owner** – User can edit, share, and delete the chat agent.
 - Change to **Viewer** – User can view and use the chat agent.
 - **Remove access** – Remove user or group access to a chat agent.

Note

You can also define access permissions to a chat agent globally, granting all users access to a chat agent and then defining more granular access permissions for specific users and groups. To do so, choose **Settings** and then use **Global settings** to turn this feature on or off.

5. Select **Share**.

View chat agents

You can view a list of the chat agents you have created or have access to. The following procedure shows how.

To view custom chat agents

1. Log in to the Amazon Quick console.
2. From the left navigation menu, select **Chat agents**
3. The chat agents you have created or have access to will be displayed on the **Chat agents** home page.

View chat agent details

You can view the details of chat agents that are either created by you or shared with you, including all spaces and actions attached to each chat agent.

To view chat agent details

1. Log in to the Amazon Quick console.
2. From the left navigation menu, select **Chat agents**
3. Then, from the list of chat agents, from the **Name** column, click on the name of the chat agent you want to view details for. A detail modal will open.

Note

You can also view chat agent details by navigating to the **Actions** column for the chat agent you want to share, selecting the menu icon, and then selecting **View chat agent details**.

4. In the chat agent details modal that opens, you will see the following details:
 - a. **Description** – The description for the chat agent.
 - b. **Instruction summary** – A summary of the instructions defined for the chat agent.
 - c. **Created by** – Details on whether the chat agent was created by you or someone else.
 - d. **Last modified** – The time the chat agent was last modified.
 - e. Under **Capabilities**:

- **Knowledge** – The knowledge sources the chat agent is either attached to or has access to. If your chat agent is linked to a space, you'll be able to see which space it is linked to.
- **Actions** – The actions that the chat agent is attached or has access to. You can also view this by clicking on **info** next to the agent name in the chat UI agent selector.

Duplicate custom chat agents

You can duplicate existing custom chat agents without building a chat agent from scratch. Users can duplicate chat agents from the list page using the **Duplicate** action button. When you duplicate a chat agent, Amazon Quick creates a new chat agent with name format "[Original name] (copy)", which you can edit. All fields and configurations from the original chat agent are copied to the new version of the chat agent. Once duplicated, you can assign custom user access permissions to the duplicated chat agent and share it.

The following procedure outlines how to duplicate chat agents.

To duplicate custom chat agents

1. Log in to the Amazon Quick console.
2. From the left navigation menu, select **Chat agents**
3. Then, from the **Actions** column for the chat agent you want to share, select the menu icon, and then select **Duplicate**.

You will be redirected to a chat agent duplication modal with prepopulated fields.

4. Make any changes you need to, including adding the right user and group permissions, and linking it to the assets you want and then select **Launch** to launch your duplicated chat agent.

Delete chat agents

You can delete a Amazon Quick chat agent you own. The following procedure shows you how.

To delete a custom chat agent

1. Log in to the Amazon Quick console.
2. From the left navigation menu, select **Chat agents**

3. Then, from the **Actions** column for the chat agent you want to share, select the menu icon, and then select **Delete**.

Chat agent context sources and best practices

When creating custom chat agents in Amazon Quick, you have multiple options to influence chat agent behavior and frame chat agent workflows. The way you configure chat agent context and instructions significantly impacts how your chat agent performs and responds to user interactions. Understanding when to use instructions, reference documents, and spaces is crucial for creating effective chat agents that meet your specific business needs and deliver consistent, reliable results. The following sections outline how to best use chat agent context and knowledge for your use case.

Topics

- [Types of agent context](#)

Types of agent context

Amazon Quick provides three distinct methods for configuring chat agent behavior, each serving different purposes in chat agent development. Understanding how each method processes content and when to use them forms the foundation of effective chat agent configuration.

The three context types available in Amazon Quick are:

- **Persona instructions** – High-level guidance that Amazon Quick enhances using generative AI
- **Reference documents** – Documents that provide exact instructions and remain active in the chat agent's memory
- **Spaces** – Collections of searchable Amazon Quick resources and documents from which chat agents retrieve information to answer questions

The following section details what these are and the best practices of using them.

Topics

- [Chat agent persona instructions](#)
- [Chat agent reference documents](#)
- [Chat agents and spaces](#)

Chat agent persona instructions

You provide chat agent persona instructions in the **Persona instructions** field when you create your custom chat agent. These persona instructions are refined using generative AI to ensure that it works with all the chat features and the resources configured for the chat agent. This approach is ideal when you want the system to build upon your foundational direction. When you use persona instructions, Amazon Quick uses generative AI to interpret and expand your persona instructions in the backend, augment simple guidance with additional context and capabilities, optimize the prompt for better chat agent performance, and enhance effectiveness through AI-driven improvements.

Use persona instructions to:

- Provide high level guidance on behavior and goals, which are enhanced by Amazon Quick AI to customize agent responses
- Instruct how to leverage reference documents for detailed interactions
- Provide general guidance on use of spaces and connectors in its responses

This approach works best for general-purpose chat agents and exploratory use cases where you want to benefit from AI assistance in chat agent design.

Note

The AI enhancement process may modify your original wording and structure to improve chat agent performance. Critical guidelines that must remain unchanged should be placed in reference documents instead.

Chat agent reference documents

Reference documents provide specific process and response templates that the chat agents need to follow. These reference documents are part of the chat agent's permanent context and go hand-in-hand with persona instructions and response style to inform its overall behavior.

Reference documents provide complete control over chat agent behavior by preserving your content exactly as written. This approach is essential when precision and consistency are paramount. When you upload reference documents, Amazon Quick preserves precise wording and

specific instructions, maintains exact formatting and structure, provides direct control over chat agent behavior, and ensures consistency across all interactions.

Use reference documents when you need:

- Exact control where precise wording and specific instructions must be preserved
- Brand consistency with specific tone, terminology, or response formats
- Complex workflows with multi-step processes that must follow exact sequences
- Predictable behavior with consistent responses across all interactions

Supported file formats:

- .pdf, .txt, .html, .md – Instructions, templates, and guidelines
- .csv – Reference data, lookup tables, and examples
- .doc, .docx – Detailed process documentation

You can attach upto 10 files of total size 50 MB as reference documents. Amazon Quick extracts text from these documents and applies a 100K character limit after processing.

File prioritization and behavior

- Behavioral frameworks and reference data from files remain active even when users attach alternative content.
- Response templates may be temporarily adjusted if users provide conflicting formatting instructions.

Reference documents serve as the chat agent's operating manual, providing the behavioral framework that remains constant across all user interactions with the chat agent.

Chat agents and spaces

Spaces provide dynamic, searchable knowledge repositories that chat agents query during conversations to retrieve specific information. When you link spaces to chat agents, Amazon Quick chat agents look for answers within the space to generate responses. Spaces excel at managing organizational knowledge that evolves over time, allowing multiple team members to contribute

and update information while maintaining appropriate security boundaries through Amazon Quick's permission system.

Use spaces with chat agents for:

- Enabling chat agents to search and retrieve current business data
- Providing chat agents access to organizational knowledge repositories
- Allowing chat agents to respect existing permission boundaries
- Connecting chat agents to collaborative document collections
- Giving chat agents access to regularly updated information
- Enabling chat agents to query live dashboards and datasets

Spaces function as the chat agent's dynamic knowledge layer, enabling real-time access to organizational information while maintaining security boundaries and collaborative workflows.

Use a chat agent

You can use a Amazon Quick chat agent to ask questions and to accomplish your tasks. When you ask a question, the chat agent analyzes the latest approved data shared with you to generate a comprehensive response.

You can ask complex questions in plain language and get a detailed response. You can also use a Quick chat agent to perform tasks for you, such as draft an email message, create a ticket, or use flows.

To use custom chat agents

1. Log in to the Quick console.
2. From the left navigation menu, select **Chat agents**
3. Then, from the **Actions** column for the chat agent you want to share, select the menu icon, and then select **Chat**.

This will open a chat interface that you use to interact with your chat agent.

You can interact with a chat agent in the following ways:

- View and access example prompts to demonstrate possible questions and tasks

- Get contextual responses from organizational data sources
- Continue or start new conversations
- Analyze data, including summarization and comparative analysis
- Take actions like scheduling meetings and creating tickets
- View source citations to responses with clickable numbered references
- Upload files
- View conversation history

** Tip**

You can also use your chat agents while automating your routine tasks with Quick Flows. For more information, see [Chat agent](#).

To learn more about chatting using chat agents, see [Using Amazon Quick chat](#).

Organize, collaborate, and share resources with spaces in Amazon Quick

A space in Amazon Quick is a collection of data and Quick resources scoped for a particular team or domain. You can use spaces to aggregate and organize files, dashboards, topics, knowledge bases, and application actions into a unified and customizable knowledge center for your team. Spaces integrate seamlessly with Quick agents for contextual conversations and are designed to scale across personal, team, and cross-team use cases.

Spaces allow your team to get the most relevant results from conversational agents and other AI tools inside Quick by grounding the results with only data relevant for your task or domain. Multiple people on the team can contribute to the knowledge inside a space; this reduces data silos and streamline information discovery. Spaces also serve as a data layer for apps in Amazon Quick applications. For more information, see [Build web applications with apps in Amazon Quick](#).

Note

Quick resources added to a space respect user access permissions. If a user can't access a resource outside a space, they won't be able to access it inside a space. Files uploaded to a space are always available to everyone with access to the space.

You can use spaces to accomplish tasks like:

- Aggregating and sharing customer feedback, allowing any team member to ask questions and get summaries
- Onboarding new team members with team documents, manuals, and processes
- Analyzing hundreds of annual reports
- Summarizing and search for action items from audio meeting recordings

Topics

- [Amazon Quick user interactions with spaces](#)
- [Amazon Quick user permission for spaces](#)
- [Understanding file upload status](#)
- [Limitations of spaces](#)

- [Creating a space](#)
- [Managing a space](#)

Amazon Quick user interactions with spaces

The following table shows how each user type in Amazon Quick interacts with spaces.

User type	Capabilities
Administrators	• Control whether users with space creation capabilities can create spaces
Author Pro and Reader Pro users	• Create and maintain spaces • Add topics, dashboards, datasets, knowledge bases, and application actions to spaces • Upload files directly into spaces • Attach an agent to a space • Interact with space data through chat • Share spaces

Amazon Quick user permission for spaces

What you can do with a space also depends on the permissions you're assigned for it. There are two permission types that users can be assigned:

- **Owner** – Owners can create, edit, share, and delete a space.
- **Viewer** – Viewers can view ask questions, and download files from spaces.

The following table outlines how user permissions determine what you can do with a Amazon Quick space:

Permissions type	Permissions
Owners	• Create a space

Permissions type	Permissions
	• Upload files to a space • Share spaces with others • Link and unlink Amazon Quick resources (topics, dashboards, datasets, knowledge bases, and application actions) to a space • Delete a space As an owner, you can designate another user co-owner of a space. If you do so, the user can undertake all post-creation management actions an owner can.
Viewers	• Download files uploaded into a space • Ask questions from data inside space • Use a specific space as context for an agent • Search for a space by name • Access a space using a direct URL • View a preset list sample questions to help get started on using the space

Understanding file upload status

When you upload files, you should see the following status messages:

Uploading

Files are being uploaded into your computer.

Processing

Files are being processed.

Text ready

The text content of the document is ready to be queried. However, the document's images and tables are still processing.

Ready

The text and media of the document has been processed.

Deleting

The document is being deleted.

Limitations of spaces

The following list outlines the limitations of spaces:

- A space can't contain other spaces.
- Users can't access Amazon Quick resources (topics, dashboards) in a space if they don't have access to the resource outside it. Sharing a space doesn't automatically give access to users to assets within it.
- If an Amazon Quick resource is linked to a space is deleted outside the space, users of the space won't get notified of the deletion. Deleted assets will only display a message indicating unavailability.
- Files uploaded to spaces will be rejected if your index data storage capacity is full. Reach out to your system administrator if this happens.
- Files uploaded to a space must be one of the following formats: .html, .xhtml, .txt, .pdf, .csv, .tsv, .xml, .json, .md, .rtf, .xslt, .eml, .vtt, .docx, .dotx, .docm, .dotm, .ppt, .pptx, .pptm, .ppsm, .ppsx, .xlsx, .xls, .xlsm, .xltx, .xltm, .c, .h, .cpp, .hpp, .cs, .java, .py, .go, .rb, .scala, .m, .sql, .clj, .mp3, .wav, .m4a, .flac, .ogg, .mp4, .mov, .m4v.
- File uploads are limited to 500 MB for PDF, Word, and PowerPoint files, and 50 MB for all other supported file types.
- Spaces support up to 10,000 files (compared to 20 files in regular chat conversations), as long as the total space storage is less than 10 GB.
- You can add at most 20 resources of a specific resource type (dashboard, topic, knowledge base, or action) to a space.

Creating a space

You can create spaces in Amazon Quick from the **Spaces** menu on the Quick home page. You can directly add files to your space or connect existing resources to it when you create it. If you upload a file that exceeds the file size limits, or is of an unsupported file type, Quick displays an error and asks you to try again. When you create a space, you are assigned as its owner by default. The following procedure outlines how to create a space.

To create a space

1. Log in to the Quick console.
2. From the left navigation menu, select **Spaces**, and then select **Create space**.
3. On the space creation page that opens, do the following:
 - a. (Optional) Enter a name for your space.
 - b. (Optional) Enter a description for your space.
 - c. Select **Add knowledge** to begin adding content to your space.

From the dropdown menu, you can select from available resource types including **Dashboards**, **Topics**, **Knowledge bases**, **Actions**, and **File uploads**.

Note

You can preview a resource before you choose to add it by clicking on its name.

After you finish adding resources, **Space knowledge** displays a list of all assets and files added to your space.

After your space is created, you can select the sparkle icon from the top left corner of the screen and start interacting with your space using any available agent.

Managing a space

You can perform the following actions for spaces.

Topics

- [Viewing a list of spaces](#)
- [Editing a space](#)
- [Adding knowledge to a space](#)
- [Viewing a list of content inside a space](#)
- [Using a space](#)
- [Sharing a space](#)
- [Deleting a space](#)
- [Removing knowledge from a space](#)
- [Adding data from external sources into your space](#)
- [Adding application actions to a space](#)

Viewing a list of spaces

You can view a list of spaces created by—or shared with—you.

To view a list of spaces

1. Log in to the Amazon Quick console.
2. From the left navigation menu, select **Spaces**.

You will see a list of spaces displayed. The list also displays the name, description, owner, and last modified date for each space. You can search for a space, and filter for recent spaces.

Editing a space

You can edit an Amazon Quick space. When you edit a space, you can perform the following actions:

- Edit space description and space name
- Add and delete any files and link or unlink any resources (dashboards, topics, knowledge bases, and application actions) that you've added.

To edit a space

1. Log in to the Amazon Quick console.

2. From the left navigation menu, select **Spaces**.
3. In **Spaces**, select the space you want to edit.
4. Inside the space, do any of the following:
 - a. To edit **Space name** or **Space description** , double click on the name or description field. Then, select the edit icon to update the description.
 - b. In the **All knowledge** section, select **Add knowledge** to add new content.
 - c. In **Amazon Quick resources** – Link and unlink resources as needed.

Adding knowledge to a space

You can add knowledge—files and Amazon Quick resources—to a space after it's been created.

To add knowledge to an existing space

1. Log in to the Amazon Quick console.
2. From the left navigation menu, select **Spaces**, and then select your space from the list of spaces displayed.
3. Select **Add knowledge** to begin adding content to your space.

From the dropdown menu, you can select from available resource types including **Dashboards**, **Topics**, **Knowledge bases**, **Actions**, and **File uploads**.

Note

You can preview a resource before you choose to add it by clicking on its name.

After you finish adding resources, **Space knowledge** displays a list of all resources and files added to your space.

Viewing a list of content inside a space

You can view a list of content inside a space.

To view a list of content within a space

1. Log in to the Amazon Quick console.

2. From the left navigation menu, select **Spaces**.
3. Select the name of the space whose content you want to view from the list of spaces.

On the space detail page, you will see the **Space knowledge** section displaying a list of files uploaded and resources linked to the space.

Using a space

After you finish creating your space and adding knowledge to it, you can start interacting with it through agents. You can interact with an agent while in a space in the following ways:

- Interact with a space using any agent

Note

When you switch between spaces or open up a space during an on-going chat Amazon Quick detects your space and focuses your agent to it.

- Interact with a specific knowledge source within a space using an agent

To use an agent to chat with a space

1. Log in to the Amazon Quick console.
2. Then, do one of the following:
 1. From the left navigation menu, select **Spaces**.
 2. In **Spaces**, select the space name.
 3. In the space that opens, start chatting with your data using the chat interface.
 1. Use the sparkle chat icon to open the chat interface.
 2. From the resource selector menu, select **All resources**, and then select **Choose a resource**.
 3. From the **Add a resource** menu, from **Spaces**, select the space you want to chat with. Then, start chatting.

Sharing a space

You can share a space with any other Amazon Quick user. When you share a space with a user, you must also choose the permission levels that user has to the space—whether owner or viewer.

To share a space

1. Log in to the Amazon Quick console.
2. From the left navigation menu, select **Spaces**, and select the name of the space you want to share. The space should open.
3. From the top right corner of your space menu, select **Share**
4. In the **Share space** modal that opens, add the users and groups you want. Choose the permissions type for each user and group you add.

Note

You can also share a space from the space home page. Locate the space you want to share, select the menu icon in the **Actions**, and then select **Share**. In the **Share space** modal that opens, add users and groups and set their permissions.

Deleting a space

You can delete a Amazon Quick space. Deleting a space deletes all files associated with the space. Deleting a space unlinks assets associated with the space without deleting the assets. You can only delete a space you own.

To delete a space

1. Log in to the Amazon Quick console.
2. From the left navigation menu, select **Spaces**, and select the name of the space you want to share. The space should open.
3. From the top right corner of your space menu, select **Delete**, and then select **Delete** from the modal that opens.

Note

You can also delete a space from the space home page. Locate the space you want to delete, select the menu icon in the **Actions**, and then select **Delete**. In the **Delete space** modal that opens, select **Delete**.

Removing knowledge from a space

You can remove knowledge (remove files and unlink assets) from a space. When you remove files, they are removed from the Amazon Quick index in which they are stored. When you remove an asset, it unlinks it from the space without deleting it from the Amazon Quick system.

To remove knowledge from an existing space

1. Log in to the Amazon Quick console.
2. From the left navigation menu, select **Spaces**, and then select your space from the list of spaces displayed.
3. In the **Space knowledge** section:
 - a. In **Files uploaded** – From the **Actions** menu select **Delete**.
 - b. In **Amazon Quick resources** – From the **Actions** menu select **Remove from this space**.

Adding data from external sources into your space

You can pull in data from Google Drive, OneDrive, Confluence, SharePoint, web pages, and Amazon S3 into your spaces to expand the context of your chat and research agents.

To add data from external sources

1. Create a space and give it a name.
2. Select the **Knowledge bases** section in the space knowledge list.
3. Select **Add knowledge bases**.
4. In **All knowledge bases** you will see application data integrations that were created for you by others in your organization. You can select one to pull into your space, or create a new one.

Note

To learn how to create a knowledge base, refer to [Working with integrations](#).

Adding application actions to a space

You can use application actions to read and write data to external SaaS applications and MCP servers from within a space. When you add actions to a space, only those actions will be considered when you select the space as your chat context.

Note

Only authors can create connectors.

To use actions in a space

1. Create a space and give it a name.
2. Select the **Actions** section in the space knowledge list.
3. Select **Add action**.
4. Select **All actions**, and choose an action that others have created in your organization.

Note

To learn how to create a connector, refer to [Working with integrations](#).

Visualize, analyze, and share data with analyses, dashboards, and reports in Amazon Quick Sight

Amazon Quick Sight is a comprehensive business intelligence service that enables you to transform raw data into meaningful insights through interactive visualizations, dashboards, and reports. Whether you're connecting to databases, preparing datasets, creating analyses, or sharing dashboards with stakeholders, Amazon Quick Sight provides the tools you need to make data-driven decisions.

Note

Amazon Quick Sight is available only for Amazon Quick accounts provisioned through the AWS Management Console. It is not available for Free or Plus accounts created at aws.com/quick.

This section covers the complete Amazon Quick Sight workflow, from initial data connection through final report sharing. You'll learn how to connect to various data sources, prepare and transform your data, create compelling visualizations, build interactive dashboards, and leverage generative BI capabilities to accelerate your analytics workflow. Each topic builds upon the previous one, providing a comprehensive guide to maximizing your use of Amazon Quick Sight's powerful features.

Topics

- [Key concepts](#)
- [Connecting to data in Amazon Quick Sight](#)
- [Refreshing data in Amazon Quick Sight](#)
- [Preparing data in Amazon Quick Sight](#)
- [Analyses and reports: Visualizing data in Amazon Quick Sight](#)
- [Sharing and subscribing to data in Amazon Quick Sight with dashboards and reports](#)
- [Exploring interactive dashboards in Amazon Quick Sight](#)
- [Gaining insights with machine learning \(ML\) in Amazon Quick Sight](#)
- [Generative BI with Quick Sight](#)
- [Troubleshooting Amazon Quick Sight](#)

- [Developing with Amazon Quick Sight](#)

Key concepts

The following terms are used throughout this section.

Data source

A connection to an external data repository such as a database, data warehouse, cloud service, or file. Data sources provide the raw data for your analyses.

Data preparation

The process of transforming data for use in an analysis. This includes filtering data, renaming fields, changing data types, adding calculated fields, and creating SQL queries to refine data.

Dataset

A prepared collection of data ready for use in analyses and dashboards. Datasets can be stored in SPICE for fast performance or queried directly from the source.

SPICE

SPICE (Super-fast, Parallel, In-memory Calculation Engine) is the in-memory engine that Quick Sight uses for high-performance analytics. SPICE speeds up analytical queries on imported data so you don't need to retrieve data from the source every time you change an analysis or update a visual.

Analysis

The workspace for creating data visualizations. Each analysis contains a collection of visualizations that you arrange and customize on one or more sheets.

Visual

A graphical representation of data such as a chart, graph, or table. All visuals begin in AutoGraph mode, which automatically selects the best visualization type for the fields you select.

Sheet

A page within an analysis that displays a set of visualizations and insights. You can add multiple sheets to an analysis and configure them to work independently or together.

Dashboard

The published version of an analysis. You share dashboards with other Quick users and control what they can do with the data.

Connecting to data in Amazon Quick Sight

People in many different roles use Amazon Quick Sight to help them do analysis and advanced calculations, design data dashboards, embed analytics, and make better-informed decisions. Before any of that can happen, someone who understands your data needs to add it to a [Quick Sight dataset](#). Quick Sight supports direct connections and uploads from a variety of [data sources](#).

Capabilities and use cases

Amazon Quick Standard edition capabilities

After your data is available in Quick Standard edition, you can do the following:

- Transform the dataset with field formatting, hierarchies, data type conversions, and calculations.
- Create one or more data analyses based on your newly created dataset.
- Share your analysis with other people so they can help design it.
- Add charts, graphs, more datasets, and multiple pages (called sheets) to your data analysis.
- Create visual appeal with customized formatting and themes.
- Make them interactive by using parameters, controls, filters, and custom actions.
- Combine data from multiple data sources, and then build new hierarchies for drilling down and calculations only available during analytics, like aggregations, window functions, and more.
- Publish your analysis as an interactive data dashboard.
- Share the dashboard so other people can use the dashboard, even if they don't use the analysis that it's based on.
- Add more data to create more analyses and dashboards.

Amazon Quick Enterprise edition capabilities

After your data is available in Quick Enterprise edition, you can do different things depending on your role. If you can build datasets, design analyses, and publish dashboards, you can do all of the things people using Standard edition can do.

In addition, these are some examples of additional tasks that you can do:

- Create analyses that use Quick Sight insights, including machine learning (ML) powered insights for forecasting, anomaly and outlier detection, and key driver identification.
- Design narrative insights with text, colors, images, and calculations.
- Add data from virtual private clouds (VPCs) and on-premises data sources, with data encryption at rest.
- Control access in datasets by adding row and column level security.
- Refresh imported datasets every hour.
- Share emailed reports.

Application development

If you develop applications or use the AWS SDKs and AWS Command Line Interface (AWS CLI), you can do the following and more:

- Add embedded analytics and embedded interactive dashboards to websites and applications.
- Use API operations to manage data sources and datasets.
- Refresh imported data more frequently by using the data ingestion API operations.
- Script, transfer, and make templates from analyses and dashboards by using API operations.
- Programmatically assign people to security roles based on settings managed by system administrators.

Administrative functions in Quick

If you perform administrative functions in Quick, you can do the following and more:

- Manage security with shared folders to organize your teams' work and help them collaborate using dashboards, analytics, and datasets.
- Add Quick Sight to your VPC to enable access to data in VPC and on-premises data sources.
- Protect sensitive data with finely grained access control to AWS data sources.
- Manually assign people to the Quick author security role so they can prepare datasets, design analytics, and publish data dashboards at a fixed cost per month.
- Manually assign people to the Quick reader security role so they can securely interact with published data dashboards on a pay-per-session basis.

Dashboard subscription

If you subscribe to dashboards, you can do the following:

- Use and subscribe to interactive dashboards designed by your team of experts.

- Enjoy a simplified uncluttered interface.
- View dashboard snapshots in email.
- Focus on making decisions with the data at your fingertips.

After you connect to or import data, you create a dataset to shape and prepare data to share and reuse. You can view your available datasets on the **Data** page, which you reach by choosing **Data** on the Amazon Quick Sight start page. You can view available data sources and create a new dataset on the **Create a Data Set** page, which you reach by choosing **Create** then **New dataset** on the **Data** page.

Topics

- [Supported data sources](#)
- [Connect to your data with integrations and datasets](#)
- [Connecting to data catalogs in Amazon Quick](#)
- [Data source quotas](#)
- [Supported data types and values](#)
- [Working with datasets](#)
- [Working with data sources in Amazon Quick Sight](#)

Supported data sources

Amazon Quick Sight supports a variety of data sources that you can use to provide data for analyses. The following data sources are supported.

Connecting to relational data

You can use any of the following relational data stores as data sources for Amazon Quick Sight:

- Amazon Athena
- Amazon Aurora
- AWS Glue Data Catalog can be accessed using AWS Glue data catalog compatible services, such as Athena or Redshift Spectrum
- Amazon OpenSearch Service
- Amazon Redshift
- Amazon Redshift Spectrum

- Amazon S3
- Amazon S3 Analytics
- Amazon S3 Tables
- Apache Impala
- Apache Spark 2.0 or later
- AWS IoT Analytics
- Databricks (E2 Platform only) on Spark 1.6 or later, up to version 3.0
- Exasol 7.1.2 or later
- Google BigQuery
- MariaDB 10.0 or later
- Microsoft SQL Server 2012 or later
- MySQL 5.7 or later

 Note

Effective October 2023, the MySQL community has deprecated support for MySQL version 5.7. This means that Amazon Quick Sight will no longer support new features, enhancements, bug fixes, or security patches for MySQL 5.7. Support for existing query workload will take place at a best effort basis. Quick Sight customers can still use MySQL 5.7 datasets with Quick Sight, but we encourage customers to upgrade their MySQL databases (DB) to major version 8.0 or higher. To see the statement provided by Amazon RDS, see [Amazon RDS Extended Support opt-in behavior is changing. Upgrade your Amazon RDS for MySQL 5.7 database instances before February 29, 2024 to avoid potential increase in charges.](#)

Amazon RDS has updated their security settings for Amazon RDS MySQL 8.3. Any connections from Quick Sight to Amazon RDS MySQL 8.3 are SSL-enabled by default. This is the only option available for MySQL 8.3. connections.

TLS 1.2 for MySQL connections requires MySQL version 5.7.28 or higher. For MySQL versions below 5.7.28, Quick Sight falls back to TLS 1.1. If your security requirements mandate TLS 1.2, ensure your MySQL or Aurora MySQL database is running version 5.7.28 or higher.

- Oracle 12c or later
- PostgreSQL 9.3.1 or later

Note

SCRAM based authentication to PostgreSQL from Amazon Quick Sight is supported for the following connectors: RDS hosted PostgreSQL, Aurora PostgreSQL, and Vanilla PostgreSQL. If the appropriate PostgreSQL engine version is used, and the correct configurations in PostgreSQL for SCRAM are set up, no additional configurations are needed in Quick Sight. If you are still experiencing issues establishing a SCRAM authentication to PostgreSQL from Quick Sight, please create a support ticket.

- Presto 0.167 or later
- Snowflake
- Starburst
- Trino
- Teradata 14.0 or later
- Timestream

Note

You can access additional data sources not listed here by linking or importing them through supported data sources.

Amazon Redshift clusters, Amazon Athena databases, and Amazon RDS instances must be in AWS. Other database instances must be in one of the following environments to be accessible from Amazon Quick Sight:

- Amazon EC2
- Local (on-premises) databases
- Data in a data center or some other internet-accessible environment

For more information, see [Infrastructure security in Amazon Quick](#).

Importing file data

You can use files in Amazon S3 or on your local (on-premises) network as data sources. Quick Sight supports files in the following formats:

- CSV and TSV – Comma-delimited and tab-delimited text files
- ELF and CLF – Extended and common log format files
- JSON – Flat or semistructured data files
- XLSX – Microsoft Excel files

Quick Sight supports UTF-8 file encoding, but not UTF-8 (with BOM).

Files in Amazon S3 that have been compressed with zip, or gzip (www.gzip.org), can be imported as-is. If you used another compression program for files in Amazon S3, or if the files are on your local network, remove compression before importing them.

JSON data

Amazon Quick Sight natively supports JSON flat files and JSON semistructured data files.

You can either upload a JSON file or connect to your Amazon S3 bucket that contains JSON data. Amazon Quick Sight automatically performs schema and type inference on JSON files and embedded JSON objects. Then it flattens the JSON, so you can analyze and visualize application-generated data.

Basic support for JSON flat-file data includes the following:

- Inferring the schema
- Determining data types
- Flattening the data
- Parsing JSON (JSON embedded objects) from flat files

Support for JSON file structures (.json) includes the following:

- JSON records with structures
- JSON records with root elements as arrays

You can also use the `parseJson` function to extract values from JSON objects in a text file. For example, if your CSV file has a JSON object embedded in one of the fields, you can extract a value from a specified key-value pair (KVP). For more information on how to do this, see [parseJson](#).

The following JSON features aren't supported:

- Reading JSON with a structure containing a list of records
- List attributes and list objects within a JSON record; these are skipped during import
- Customizing upload or configuration settings
- `parseJSON` functions for SQL and analyses
- Error messaging for invalid JSON
- Extracting a JSON object from a JSON structure
- Reading delimited JSON records

You can use the `parseJson` function to parse flat files during data preparation. This function extracts elements from valid JSON structures and lists.

The following JSON values are supported:

- JSON object
- String (double quoted)
- Number (integer and float)
- Boolean
- NULL

Software as a service (SaaS) data

Quick Sight can connect to a variety of Software as a Service (SaaS) data sources either by connecting directly or by using Open Authorization (OAuth).

SaaS sources that support direct connection include the following:

- Jira
- ServiceNow

SaaS sources that use OAuth require that you authorize the connection on the SaaS website. For this to work, Quick Sight must be able to access the SaaS data source over the network. These sources include the following:

- Adobe Analytics
- GitHub
- Salesforce

You can use reports or objects in the following editions of Salesforce as data sources for Amazon Quick Sight:

- Enterprise Edition
- Unlimited Edition
- Developer Edition

Local data sources

To connect to on premises data sources, you need to add your data sources and a Quick-specific network interface to Amazon Virtual Private Cloud (Amazon VPC). When configured properly, a VPC based on Amazon VPC resembles a traditional network that you operate in your own data center. It enables you to secure and isolate traffic between resources. You define and control the network elements to suit your requirements, while still getting the benefit of cloud networking and the scalable infrastructure of AWS.

For detailed information, see [Infrastructure security in Amazon Quick](#).

Connect to your data with integrations and datasets

You can connect Amazon Quick Sight to different types of data sources. This includes data residing in Software-as-a-Service (SaaS) applications, flat files stored in Amazon S3 buckets, data from third-party services like Salesforce, and query results from Athena. Use the following examples to learn more about the requirements for connecting to specific data sources.

Topics

- [Creating a dataset using Amazon Athena data](#)
- [Using Amazon OpenSearch Service with Amazon Quick Sight](#)
- [Creating a dataset using Amazon S3 files](#)
- [Creating a dataset using Amazon S3 Tables](#)

- [Creating a data source using Apache Spark](#)
- [Using Databricks in Quick Sight](#)
- [Creating a dataset using Google BigQuery](#)
- [Creating a dataset using a Google Sheets data source](#)
- [Creating a dataset using an Apache Impala data source](#)
- [Creating a dataset using a Microsoft Excel file](#)
- [Creating a data source using Presto](#)
- [Using Snowflake with Amazon Quick Sight](#)
- [Using Starburst with Amazon Quick Sight](#)
- [Creating a data source and data set from SaaS sources](#)
- [Creating a dataset from Salesforce](#)
- [Using Trino with Amazon Quick Sight](#)
- [Creating a dataset using a local text file](#)
- [Using Amazon Timestream data with Amazon Quick Sight](#)

Creating a dataset using Amazon Athena data

Use the following procedure to create a new dataset that connects to Amazon Athena data or to Athena Federated Query data.

To connect to Amazon Athena

1. Begin by creating a new dataset. Choose **Data** from the navigation pane at left.
2. Choose **Create**, then choose **New dataset**.
3. a. To use an existing Athena connection profile (common), choose the card for the existing data source that you want to use. Choose **Select**.

Cards are labeled with the Athena data source icon and the name provided by the person who created the connection.

- b. To create a new Athena connection profile (less common), use the following steps:
 1. Choose **New data source**, then choose the **Athena** data source card.
 2. Choose **Next**.
 3. For **Data source name**, enter a descriptive name.

4. For **Athena workgroup**, choose your workgroup.
 5. Choose **Validate connection** to test the connection.
 6. Choose **Create data source**.
 7. (Optional) Select an IAM role ARN for queries to run as.
4. On the **Choose your table** screen, do the following:
- a. For **Catalog**, choose one of the following:
 - If you are using Athena Federated Query, choose the catalog you want to use.
 - Otherwise, choose **AwsDataCatalog**.
 - b. Choose one of the following:
 - To write a SQL query, choose **Use custom SQL**.
 - To choose a database and table, choose your catalog that contains your databases from the dropdown under **Catalog**. Then, choose a database from the dropdown under **Database** and choose a table from the **Tables** list that appears for your database.
- If you don't have the right permissions, you receive the following error message: "You don't have sufficient permissions to connect to this dataset or run this query." Contact your Quick administrator for assistance. For more information, see [Authorizing connections to Amazon Athena](#).
5. Choose **Edit/preview data**.
 6. Create a dataset and analyze the data using the table by choosing **Visualize**. For more information, see [Analyses and reports: Visualizing data in Amazon Quick Sight](#).

Using Amazon OpenSearch Service with Amazon Quick Sight

Following, you can find how to connect to your Amazon OpenSearch Service data using Amazon Quick Sight.

Creating a new Quick Sight data source connection for OpenSearch Service

Following, you can find how to connect to OpenSearch Service

Before you can proceed, Amazon Quick Sight needs to be authorized to connect to Amazon OpenSearch Service. If connections aren't enabled, you get an error when you try to connect. A Quick Sight administrator can authorize connections to AWS resources.

To authorize Quick Sight to initiate a connection to OpenSearch Service

1. Open the menu by clicking on your profile icon at top right, then choose **Manage Quick**. If you don't see the **Manage Quick** option on your profile menu, ask your Amazon Quick administrator for assistance.
2. Choose **Security & permissions, Add or remove**.
3. Enable the option for **OpenSearch**.
4. Choose **Update**.

After OpenSearch Service is accessible, you create a data source so people can use the specified domains.

To connect to OpenSearch Service

1. Begin by creating a new dataset. Choose **Data** from the navigation pane at left, then choose **Create and New Dataset**.
2. Choose the **Amazon OpenSearch** data source card.
3. For **Data source name**, enter a descriptive name for your OpenSearch Service data source connection, for example OpenSearch Service ML Data. Because you can create many datasets from a connection to OpenSearch Service, it's best to keep the name simple.
4. For **Connection type**, choose the network you want to use. This can be a virtual private cloud (VPC) based on Amazon VPC or a public network. The list of VPCs contains the names of VPC connections, rather than VPC IDs. These names are defined by the Quick administrator.
5. For **Domain**, choose the OpenSearch Service domain that you want to connect to.
6. Choose **Validate connection** to check that you can successfully connect to OpenSearch Service.
7. Choose **Create data source** to proceed.
8. For **Tables**, choose the one you want to use, then choose **Select** to continue.
9. Do one of the following:
 - To import your data into the Quick Sight in-memory engine (called SPICE), choose **Import to SPICE for quicker analytics**. For information about how to enable importing OpenSearch data, see [Authorizing connections to Amazon OpenSearch Service](#).
 - To allow Quick Sight to run a query against your data each time you refresh the dataset or use the analysis or dashboard, choose **Directly query your data**.

To enable autorefresh on a published dashboard that uses OpenSearch Service data, the OpenSearch Service dataset needs to use a direct query.

10. Choose **Edit/Preview** and then **Save** to save your dataset and close it.

Managing permissions for OpenSearch Service data

The following procedure describes how to view, add, and revoke permissions to allow access to the same OpenSearch Service data source. The people that you add need to be active users in Quick Sight before you can add them.

To edit permissions on a data source

1. Choose **Data** at left, then scroll down to find the data source card for your Amazon OpenSearch Service connection. An example might be `US Amazon OpenSearch Service Data`.
2. Choose the **Amazon OpenSearch** dataset.
3. On the dataset details page that opens, choose the **Permissions** tab.

A list of current permissions appears.

4. To add permissions, choose **Add users & groups**, then follow these steps:
 - a. Add users or groups to allow them to use the same dataset.
 - b. When you're finished adding everyone that you want to add, choose the **Permissions** that you want to apply to them.
5. (Optional) To edit permissions, you can choose **Viewer** or **Owner**.
 - Choose **Viewer** to allow read access.
 - Choose **Owner** to allow that user to edit, share, or delete this Quick Sight dataset.
6. (Optional) To revoke permissions, choose **Revoke access**. After you revoke someone's access, they can't create new datasets from this data source. However, their existing datasets still have access to this data source.
7. When you are finished, choose **Close**.

Adding a new Quick Sight dataset for OpenSearch Service

After you have an existing data source connection for OpenSearch Service, you can create OpenSearch Service datasets to use for analysis.

To create a dataset using OpenSearch Service

1. From the start page, choose **Data, Create, New dataset**.
2. Scroll down to the data source card for your OpenSearch Service connection. If you have many data sources, you can use the search bar at the top of the page to find your data source with a partial match on the name.
3. Choose the **Amazon OpenSearch** data source card, and then choose **Create data set**.
4. For **Tables**, choose the OpenSearch Service index that you want to use.
5. Choose **Edit/Preview**.
6. Choose **Save** to save and close the dataset.

Adding OpenSearch Service data to an analysis

After you have an OpenSearch Service dataset available, you can add it to a Quick Sight analysis. Before you begin, make sure that you have an existing dataset that contains the OpenSearch Service data that you want to use.

To add OpenSearch Service data to an analysis

1. Choose **Analyses** at left.
2. Do one of the following:
 - To create a new analysis, choose **New analysis** at right.
 - To add to an existing analysis, open the analysis that you want to edit.
 - Choose the pencil icon near at top left.
 - Choose **Add data set**.
3. Choose the OpenSearch Service dataset that you want to add.

For information on using OpenSearch Service in visualizations, see [Limitations for using OpenSearch Service](#).

4. For more information, see [Working with analyses](#).

Limitations for using OpenSearch Service

The following limitations apply to using OpenSearch Service datasets:

- OpenSearch Service datasets support a subset of the visual types, sort options, and filter options.
- To enable autorefresh on a published dashboard that uses OpenSearch Service data, the OpenSearch Service dataset needs to use a direct query.
- Multiple subquery operations aren't supported. To avoid errors during visualization, don't add multiple fields to a field well, use one or two fields per visualization, and avoid using the **Color** field well.
- Custom SQL isn't supported.
- Crossdataset joins and self joins aren't supported.
- Calculated fields aren't supported.
- Text fields aren't supported.
- The "other" category isn't supported. If you use an OpenSearch Service dataset with a visualization that supports the "other" category, disable the "other" category by using the menu on the visual.

Creating a dataset using Amazon S3 files

To create a dataset using one or more text files (.csv, .tsv, .clf, or .elf) from Amazon S3, create a manifest for Quick Sight. Quick Sight uses this manifest to identify the files that you want to use and to the upload settings needed to import them. When you create a dataset using Amazon S3, the file data is automatically imported into [SPICE](#).

You must grant Quick Sight access to any Amazon S3 buckets that you want to read files from. For information about granting Quick Sight access to AWS resources, see [Configuring Amazon Quick Sight access to AWS data sources](#).

Topics

- [Supported formats for Amazon S3 manifest files](#)
- [Creating Amazon S3 datasets](#)
- [Datasets using S3 files in another AWS account](#)

Supported formats for Amazon S3 manifest files

You use JSON manifest files to specify files in Amazon S3 to import into Quick Sight. These JSON manifest files can use either the Quick Sight format described following or the Amazon Redshift format described in [Using a manifest to specify data files](#) in the *Amazon Redshift Database Developer Guide*. You don't have to use Amazon Redshift to use the Amazon Redshift manifest file format.

If you use an Quick Sight manifest file, it must have a .json extension, for example `my_manifest.json`. If you use an Amazon Redshift manifest file, it can have any extension.

If you use an Amazon Redshift manifest file, Quick Sight processes the optional mandatory option as Amazon Redshift does. If the associated file isn't found, Quick Sight ends the import process and returns an error.

Files that you select for import must be delimited text (for example, .csv or .tsv), log (.clf), or extended log (.elf) format, or JSON (.json). All files identified in one manifest file must use the same file format. Plus, they must have the same number and type of columns. Quick Sight supports UTF-8 file encoding, but not UTF-8 with byte-order mark (BOM). If you are importing JSON files, then for `globalUploadSettings` specify `format`, but not `delimiter`, `textqualifier`, or `containsHeader`.

Make sure that any files that you specify are in Amazon S3 buckets that you have granted Quick Sight access to. For information about granting Quick Sight access to AWS resources, see [Configuring Amazon Quick Sight access to AWS data sources](#).

Manifest file format for Quick Sight

Quick Sight manifest files use the following JSON format.

```
{
  "fileLocations": [
    {
      "URIs": [
        "uri1",
        "uri2",
        "uri3"
      ]
    },
    {
      "URIPrefixes": [
```



```
        "prefix1",
        "prefix2",
        "prefix3"
    ]
  },
  "globalUploadSettings": {
    "format": "JSON",
    "delimiter": ",",
    "textqualifier": "'",
    "containsHeader": "true"
  }
}
```

Use the fields in the `fileLocations` element to specify the files to import, and the fields in the `globalUploadSettings` element to specify import settings for those files, such as field delimiters.

The manifest file elements are described following:

- **fileLocations** – Use this element to specify the files to import. You can use either or both of the URIs and `URIPrefixes` arrays to do this. You must specify at least one value in one or the other of them.
- **URIs** – Use this array to list URIs for specific files to import.

Quick Sight can access Amazon S3 files that are in any AWS Region. However, you must use a URI format that identifies the AWS Region of the Amazon S3 bucket if it's different from that used by your Quick account.

URIs in the following formats are supported.

URI format	Example	Comments
<code>https://s3.amazonaws.com/<bucket name>/<file name></code>	<code>https://s3.amazonaws.com/amzn-s3-demo-bucket/data.csv</code>	
<code>s3://<bucket name>/<file name></code>	<code>s3://amzn-s3-demo-bucket/data.csv</code>	

URI format	Example	Comments
<code>https://<bucket name>.s3.amazonaws.com/<file name></code>	<code>https://<i>amzn-s3-demo-bucket</i>.s3.amazonaws.com/<i>data.csv</i></code>	
<code>https://s3-<region name>.amazonaws.com/<bucket name>/<file name></code>	<code>https://<i>s3-us-east-1.amazonaws.com</i> /<i>amzn-s3-demo-bucket</i> /<i>data.csv</i></code>	This URI type identifies the AWS Region for the Amazon S3 bucket.
<code>https://<bucket name>.s3-<region name>.amazonaws.com/<file name></code>	<code>https://<i>amzn-s3-demo-bucket</i>.s3-us-east-1.amazonaws.com/<i>data.csv</i></code>	This URI type identifies the AWS Region for the Amazon S3 bucket.

- **URIPrefixes** – Use this array to list URI prefixes for S3 buckets and folders. All files in a specified bucket or folder are imported. Quick Sight recursively retrieves files from child folders.

Quick Sight can access Amazon S3 buckets or folders that are in any AWS Region. Make sure to use a URI prefix format that identifies the S3 bucket's AWS Region if it's different from that used by your Quick account.

URI prefixes in the following formats are supported.

URI prefix format	Example	Comments
<code>https://s3.amazonaws.com/<bucket name>/</code>	<code>https://s3.amazonaws.com/amzn-s3-demo-bucket/</code>	
<code>https://s3.amazonaws.com/<bucket name>/<folder name1>/(<folder name2>/etc.)</code>	<code>https://s3.amazonaws.com/amzn-s3-demo-bucket/folder1/</code>	
<code>s3://<bucket name></code>	<code>s3://amzn-s3-demo-bucket/</code>	

URI prefix format	Example	Comments
<code>s3://<bucket name>/<folder name1>/(<folder name2>/etc.)</code>	<code>s3://amzn-s3-demo-bucket/folder1/</code>	
<code>https://<bucket name>.s3.amazonaws.com</code>	<code>https://amzn-s3-demo-bucket.s3.amazonaws.com</code>	
<code>https://s3-<region name>.amazonaws.com/<bucket name>/</code>	<code>https://s3-your-region-for-example-us-east-2.amazonaws.com/amzn-s3-demo-bucket /</code>	This URIPrefix type identifies the AWS Region for the Amazon S3 bucket.
<code>https://s3-<region name>.amazonaws.com/<bucket name>/<folder name1>/(<folder name2>/etc.)</code>	<code>https://s3-us-east-1.amazonaws.com/amzn-s3-demo-bucket /folder1/</code>	This URIPrefix type identifies the AWS Region for the Amazon S3 bucket.
<code>https://<bucket name>.s3-<region name>.amazonaws.com</code>	<code>https://amzn-s3-demo-bucket.s3-us-east-1.amazonaws.com</code>	This URIPrefix type identifies the AWS Region for the Amazon S3 bucket.

- **globalUploadSettings** – (Optional) Use this element to specify import settings for the Amazon S3 files, such as field delimiters. If this element is not specified, Quick Sight uses the default values for the fields in this section.

⚠ Important

For log (.clf) and extended log (.elf) files, only the **format** field in this section is applicable, so you can skip the other fields. If you choose to include them, their values are ignored.

- **format** – (Optional) Specify the format of the files to be imported. Valid formats are **CSV**, **TSV**, **CLF**, **ELF**, and **JSON**. The default value is **CSV**.
- **delimiter** – (Optional) Specify the file field delimiter. This must map to the file type specified in the **format** field. Valid formats are commas (,) for .csv files and tabs (\t) for .tsv files. The default value is comma (,).
- **textqualifier** – (Optional) Specify the file text qualifier. Valid formats are single quote ('), double quotes (\"). The leading backslash is a required escape character for a double quote in JSON. The default value is double quotes (\"). If your text doesn't need a text qualifier, don't include this property.
- **containsHeader** – (Optional) Specify whether the file has a header row. Valid formats are **true** or **false**. The default value is **true**.

Manifest file examples for Quick Sight

The following are some examples of completed Quick Sight manifest files.

The following example shows a manifest file that identifies two specific .csv files for import. These files use double quotes for text qualifiers. The **format**, **delimiter**, and **containsHeader** fields are skipped because the default values are acceptable.

```
{
  "fileLocations": [
    {
      "URIs": [
        "https://yourBucket.s3.amazonaws.com/data-file.csv",
        "https://yourBucket.s3.amazonaws.com/data-file-2.csv"
      ]
    }
  ],
  "globalUploadSettings": {
    "textqualifier": "\""
  }
}
```

```
}
}
```

The following example shows a manifest file that identifies one specific .tsv file for import. This file also includes a bucket in another AWS Region that contains additional .tsv files for import. The `textqualifier` and `containsHeader` fields are skipped because the default values are acceptable.

```
{
  "fileLocations": [
    {
      "URIs": [
        "https://s3.amazonaws.com/amzn-s3-demo-bucket/data.tsv"
      ]
    },
    {
      "URIPrefixes": [
        "https://s3-us-east-1.amazonaws.com/amzn-s3-demo-bucket/"
      ]
    }
  ],
  "globalUploadSettings": {
    "format": "TSV",
    "delimiter": "\t"
  }
}
```

The following example identifies two buckets that contain .clf files for import. One is in the same AWS Region as the Quick account, and one in a different AWS Region. The `delimiter`, `textqualifier`, and `containsHeader` fields are skipped because they are not applicable to log files.

```
{
  "fileLocations": [
    {
      "URIPrefixes": [
        "https://amzn-s3-demo-bucket1.your-s3-url.com",
        "s3://amzn-s3-demo-bucket2/"
      ]
    }
  ],
  "globalUploadSettings": {
```

```
    "format": "CLF"
  }
}
```

The following example uses the Amazon Redshift format to identify a .csv file for import.

```
{
  "entries": [
    {
      "url": "https://amzn-s3-demo-bucket.your-s3-url.com/myalias-test/file-to-import.csv",
      "mandatory": true
    }
  ]
}
```

The following example uses the Amazon Redshift format to identify two JSON files for import.

```
{
  "fileLocations": [
    {
      "URIs": [
        "https://yourBucket.s3.amazonaws.com/data-file.json",
        "https://yourBucket.s3.amazonaws.com/data-file-2.json"
      ]
    }
  ],
  "globalUploadSettings": {
    "format": "JSON"
  }
}
```

Creating Amazon S3 datasets

To create an Amazon S3 dataset

1. Check [Data source quotas](#) to make sure that your target file set doesn't exceed data source quotas.
2. Create a manifest file to identify the text files that you want to import, using one of the formats specified in [Supported formats for Amazon S3 manifest files](#).
3. Save the manifest file to a local directory, or upload it into Amazon S3.

4. On the Quick start page, choose **Data**.
5. On the **Data** page, choose **Create** then **New dataset**.
6. Choose the Amazon S3 icon and then choose **Next**.
7. For **Data source name**, enter a description of the data source. This name should be something that helps you distinguish this data source from others.
8. For **Upload a manifest file**, do one of the following:
 - To use a local manifest file, choose **Upload**, and then choose **Upload a JSON manifest file**. For **Open**, choose a file, and then choose **Open**.
 - To use a manifest file from Amazon S3, choose **URL**, and enter the URL for the manifest file. To find the URL of a pre-existing manifest file in the Amazon S3 console, navigate to the appropriate file and choose it. A properties panel displays, including the link URL. You can copy the URL and paste it into Quick Sight.
9. Choose **Connect**.
10. To make sure that the connection is complete, choose **Edit/Preview data**. Otherwise, choose **Visualize** to create an analysis using the data as-is.

If you choose **Edit/Preview data**, you can specify a dataset name as part of preparing the data. Otherwise, the dataset name matches the name of the manifest file.

To learn more about data preparation, see [Preparing data in Amazon Quick Sight](#).

Creating datasets based on multiple Amazon S3 files

You can use one of several methods to merge or combine files from Amazon S3 buckets inside Quick Sight:

- **Combine files by using a manifest** – In this case, the files must have the same number of fields (columns). The data types must match between fields in the same position in the file. For example, the first field must have the same data type in each file. The same goes for the second field, and the third field, and so on. Quick Sight takes field names from the first file.

The files must be listed explicitly in the manifest. However, they don't have to be inside the same Amazon S3 bucket.

In addition, the files must follow the rules described in [Supported formats for Amazon S3 manifest files](#).

For more details about combining files using a manifest, see [Creating a dataset using Amazon S3 files](#).

- **Merge files without using a manifest** – To merge multiple files into one without having to list them individually in the manifest, you can use Athena. With this method, you can simply query your text files, like they are in a table in a database. For more information, see the post [Analyzing data in Amazon S3 using Athena](#) in the Big Data blog.
- **Use a script to append files before importing** – You can use a script designed to combine your files before uploading.

Datasets using S3 files in another AWS account

Use this section to learn how to set up security so you can use Quick Sight to access Amazon S3 files in another AWS account.

For you to access files in another account, the owner of the other account must first set Amazon S3 to grant you permissions to read the file. Then, in Quick Sight, you must set up access to the buckets that were shared with you. After both of these steps are finished, you can use a manifest to create a dataset.

Note

To access files that are shared with the public, you don't need to set up any special security. However, you still need a manifest file.

Topics

- [Setting up Amazon S3 to allow access from a different Quick account](#)
- [Setting up Quick Sight to access Amazon S3 files in another AWS account](#)

Setting up Amazon S3 to allow access from a different Quick account

Use this section to learn how to set permissions in Amazon S3 files so they can be accessed by Quick Sight in another AWS account.

For information on accessing another account's Amazon S3 files from your Quick Sight account, see [Setting up Quick Sight to access Amazon S3 files in another AWS account](#). For more information

about S3 permissions, see [Managing access permissions to your Amazon S3 resources](#) and [How do I set permissions on an object?](#)

You can use the following procedure to set this access from the S3 console. Or you can grant permissions by using the AWS CLI or by writing a script. If you have a lot of files to share, you can instead create an S3 bucket policy on the `s3:GetObject` action. To use a bucket policy, add it to the bucket permissions, not to the file permissions. For information on bucket policies, see [Bucket policy examples](#) in the *Amazon S3 Developer Guide*.

To set access from a different Quick account from the S3 console

1. Get the email address of the AWS account email that you want to share with. Or you can get and use the canonical user ID. For more information on canonical user IDs, see [AWS account identifiers](#) in the *AWS General Reference*.
2. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
3. Find the Amazon S3 bucket that you want to share with Quick Sight. Choose **Permissions**.
4. Choose **Add Account**, and then enter an email address, or paste in a canonical user ID, for the AWS account that you want to share with. This email address should be the primary one for the AWS account.
5. Choose **Yes** for both **Read bucket permissions** and **List objects**.

Choose **Save** to confirm.

6. Find the file that you want to share, and open the file's permission settings.
7. Enter an email address or the canonical user ID for the AWS account that you want to share with. This email address should be the primary one for the AWS account.
8. Enable **Read object** permissions for each file that Quick Sight needs access to.
9. Notify the Quick user that the files are now available for use.

Setting up Quick Sight to access Amazon S3 files in another AWS account

Use this section to learn how to set up Quick Sight so you can access Amazon S3 files in another AWS account. For information on allowing someone else to access your Amazon S3 files from their Quick account, see [Setting up Amazon S3 to allow access from a different Quick account](#).

Use the following procedure to access another account's Amazon S3 files from Quick Sight. Before you can use this procedure, the users in the other AWS account must share the files in their Amazon S3 bucket with you.

To access another account's Amazon S3 files from Quick Sight

1. Verify that the user or users in the other AWS account gave your account read and write permission to the S3 bucket in question.
2. Choose your profile icon, and then choose **Manage Quick Sight**.
3. Choose **Security & permissions**.
4. Under **Quick Sight access to AWS services**, choose **Manage**.
5. Choose **Select S3 buckets**.
6. On the **Select Amazon S3 buckets** screen, choose the **S3 buckets you can access across AWS** tab.

The default tab is named **S3 buckets linked to Quick Sight account**. It shows all the buckets your Quick account has access to.

7. Do one of the following:
 - To add all the buckets that you have permission to use, choose **Choose accessible buckets from other AWS accounts**.
 - If you have one or more Amazon S3 buckets that you want to add, enter their names. Each must exactly match the unique name of the Amazon S3 bucket.

If you don't have the appropriate permissions, you see the error message "We can't connect to this S3 bucket. Make sure that any S3 buckets you specify are associated with the AWS account used to create this Quick account." This error message appears if you don't have either account permissions or Quick Sight permissions.

Note

To use Amazon Athena, Quick Sight needs to access the Amazon S3 buckets that Athena uses.

You can add them here one by one, or use the **Choose accessible buckets from other AWS accounts** option.

8. Choose **Select buckets** to confirm your selection.

9. Create a new dataset based on Amazon S3, and upload your manifest file. For more information Amazon S3 datasets, see [Creating a dataset using Amazon S3 files](#).

Creating a dataset using Amazon S3 Tables

You can use Amazon S3 Tables as a data source to create datasets in Quick Sight. Amazon S3 Tables stores data as Apache Iceberg tables, enabling you to query your data lake directly from Quick Sight without setting up an intermediate query layer.

You can use Amazon S3 Tables datasets in the following query modes:

- **Direct Query** – Query your Amazon S3 Tables data directly without importing it.
- **SPICE** – Import your Amazon S3 Tables data into SPICE for faster performance.

Configuring Amazon S3 Tables resource permissions (admin)

Before authors can create datasets from Amazon S3 Tables, an admin must configure resource permissions.

To configure Amazon S3 Tables resource permissions

1. Log in as an admin user.
2. Choose the people icon in the top right, and then choose **Manage QuickSuite** to open the admin page.
3. Choose the **AWS Resources** card to open the resource permissions page.
4. Choose **Use QuickSuite-managed role (default)** so that Quick manages the related IAM role and permissions for you.
5. Select **Amazon S3 Tables**. All available S3 table buckets in your current Region are displayed.
6. Select the S3 table buckets that you want to allow authors to use, and then choose **Finish**.
7. Choose **Save**.

After you save, a new IAM role named `aws-quicksight-s3-tables-role-v0` is created with permissions to read the selected S3 table buckets.

Creating a data source for an Amazon S3 table bucket

To create an Amazon S3 Tables data source

1. From the left navigation pane, choose **Data**, and then choose the **Data sources** tab.
2. Choose **Create data source**.
3. Choose the **Amazon S3 Tables** card.
4. Enter the ARN of the S3 table bucket and a name for the data source, and then choose **Create data source**.

After a few seconds, the data source appears in your data sources list.

Creating a dataset from an Amazon S3 Tables data source

To create a dataset using Amazon S3 Tables

1. From the left navigation pane, choose **Data**, and then choose the **Datasets** tab.
2. Choose **Create dataset**.
3. Select the Amazon S3 Tables data source that you want to use.
4. Select the namespace of the S3 table that you want to use, and then select the table.
5. Choose **Direct Query** or **SPICE** mode based on your use case.
6. Choose **Edit/Preview data** to transform your data in the same way you would for other data sources. After all transformations are defined, choose **Save and publish** to save the dataset.

Note

Amazon S3 Tables datasets support a maximum of 2 billion rows.

For information about supported AWS Regions, see [Supported AWS Regions](#).

Creating a data source using Apache Spark

You can connect directly to Apache Spark using Quick Sight, or you can connect to Spark through Spark SQL. Using the results of queries, or direct links to tables or views, you create data sources in Quick Sight. You can either directly query your data through Spark, or you can import the results of your query into [SPICE](#).

Before you use Quick Sight with Spark products, you must configure Spark for Quick Sight.

Quick Sight requires your Spark server to be secured and authenticated using LDAP, which is available to Spark version 2.0 or later. If Spark is configured to allow unauthenticated access, Quick Sight refuses the connection to the server. To use Quick Sight as a Spark client, you must configure LDAP authentication to work with Spark.

The Spark documentation contains information on how to set this up. To start, you need to configure it to enable front-end LDAP authentication over HTTPS. For general information on Spark, see [the Apache spark website](#). For information specifically on Spark and security, see [Spark security documentation](#).

To make sure that you have configured your server for Quick Sight access, follow the instructions in [Network and database configuration requirements](#).

Using Databricks in Quick Sight

Use this section to learn how to connect from Quick Sight to Databricks.

To connect to Databricks

1. Begin by creating a new dataset. Choose **Data** from the navigation pane at left.
2. Choose **Create** then **New Dataset**.
3. Choose the **Databricks** data source card.
4. For **Data source name**, enter a descriptive name for your Databricks data source connection, for example Databricks CS. Because you can create many datasets from a connection to Databricks, it's best to keep the name simple.
5. For **Connection type**, select the type of network you're using.
 - **Public network** – if your data is shared publicly.
 - **VPC** – if your data is inside a VPC.

Note

If you're using VPC, and you don't see it listed, check with your administrator.

6. For **Database server**, enter the **Hostname of workspace** specified in your Databricks connection details.

7. For **HTTP Path**, enter the **Partial URL for the spark instance** specified in your Databricks connection details.
8. For **Port**, enter the **port** specified in your Databricks connection details.
9. For **Username** and **Password**, enter your connection credentials.
10. To verify the connection is working, click **Validate connection**.
11. To finish and create the data source, click **Create data source**.

Adding a new Quick Sight dataset for Databricks

After you have an existing data source connection for Databricks data, you can create Databricks datasets to use for analysis.

To create a dataset using Databricks

1. Choose **Data** at left, then scroll down to find the data source card for your Databricks connection. If you have many data sources, you can use the search bar at the top of the page to find your data source with a partial match on the name.
2. Choose the **Databricks** data source card, and then choose **Create data set**.
3. To specify the table you want to connect to, first select the Catalog and Schema you want to use. Then, for **Tables**, select the table that you want to use. If you prefer to use your own SQL statement, select **Use custom SQL**.
4. Choose **Edit/Preview**.
5. (Optional) To add more data, use the following steps:
 - a. Choose **Add data** at top right.
 - b. To connect to different data, choose **Switch data source**, and choose a different dataset.
 - c. Follow the UI prompts to finish adding data.
 - d. After adding new data to the same dataset, choose **Configure this join** (the two red dots). Set up a join for each additional table.
 - e. If you want to add calculated fields, choose **Add calculated field**.
 - f. To add a model from SageMaker AI, choose **Augment with SageMaker**. This option is only available in Quick Enterprise edition.
 - g. Clear the check box for any fields that you want to omit.
 - h. Update any data types that you want to change.

6. When you are done, choose **Save** to save and close the dataset.

Quick Sight Administrator's guide to connecting Databricks

You can use Amazon Quick Sight to connect to Databricks on AWS. You can connect to Databricks on AWS whether you signed up for through AWS Marketplace or through the Databricks website.

Before you can connect to Databricks, you create or identify existing resources that the connection requires. Use this section to help you gather the resources you need to connect from Quick Sight to Databricks.

- To learn how to obtain your Databricks connection details, see [Databricks ODBC and JDBC connections](#).
- To learn how to obtain your Databricks credentials—personal access token or user name and password—for authentication, see [Authentication requirements](#) in the [Databricks documentation](#).

To connect to a Databricks cluster, you need `Can Attach To` and `Can Restart` permissions. These permissions are managed in Databricks. For more information, see [Permission Requirements](#) in the [Databricks documentation](#).

- If you are setting up a private connection for Databricks, you can learn more about how to configure a VPC for use with Quick Sight, see [Connecting to a VPC with Amazon Quick Sight](#) in the Quick Sight documentation. If the connection isn't visible, verify with a system administrator that the network has open [inbound endpoints for Amazon Route 53](#). The hostname of a Databricks workspace uses a public IP, there needs to be DNS TCP and DNS UDP inbound and outbound rules to allow traffic on DNS port 53, for the Route 53 security group. An administrator needs to create a security group with 2 inbound rules: one for DNS(TCP) on port 53 to the VPC CIDR and one for DNS(UDP) for port 53 to the VPC CIDR.

For Databricks-related details if you are using PrivateLink instead of a public connection, see [Enable AWS PrivateLink](#) in the [Databricks documentation](#).

Creating a dataset using Google BigQuery

Note

When Quick Sight uses and transfers information that is received from Google APIs, it adheres to the [Google API Services User Data Policy](#).

Google BigQuery is a fully managed serverless data warehouse that customers use to manage and analyze their data. Google BigQuery customers use SQL to query their data without any infrastructure management.

Creating a data source connection with Google BigQuery

Prerequisites

Before you start, make sure that you have the following. These are all required to create a data source connection with Google BigQuery:

- **Project ID** – The project ID that is associated with your Google account. To find this, navigate to the Google Cloud console and choose the name of the project that you want to connect to Quick Sight. Copy the project ID that appears in the new window and record it for later use.
- **Dataset Region** – The Google region that the Google BigQuery project exists in. To find the dataset region, navigate to the Google BigQuery console and choose **Explorer**. Locate and expand the project that you want to connect to, then choose the dataset that you want to use. The dataset region appears in the pop-up that opens.
- **Google account login credentials** – The login credentials for your Google account. If you don't have this information, contact your Google account administrator.
- **Google BigQuery Permissions** – To connect your Google account with Quick Sight, make sure that your Google account has the following permissions:
 - BigQuery Job User at the Project level.
 - BigQuery Data Viewer at the Dataset or Table level.
 - BigQuery Metadata Viewer at the Project level.

For information about how to retrieve the previous prerequisite information, see [Unlock the power of unified business intelligence with Google Cloud BigQuery and Quick Sight](#).

Use the following procedure to connect your Quick account with your Google BigQuery data source.

To create a new connection to a Google BigQuery data source from Quick Sight

1. Open the [Quick console](#).
2. From the left navigation pane, choose **Data**.
3. Choose **Create** then choose **New Dataset**
4. Choose the **Google BigQuery** tile.
5. Add the data source details that you recorded in the prerequisites section earlier:
 - **Data source name** – A name for the data source.
 - **Project ID** – A Google Platform project ID. This field is case sensitive.
 - **Dataset Region** – The Google cloud platform dataset region of the project that you want to connect to.
6. Choose **Sign in**.
7. In the new window that opens, enter the login credentials for the Google account that you want to connect to.
8. Choose **Continue** to grant Quick Sight access to Google BigQuery.
9. After you create the new data source connection, continue to [Step 4](#) in the following procedure.

Adding a new Quick Sight dataset for Google BigQuery

After you create a data source connection with Google BigQuery, you can create Google BigQuery datasets for analysis. Datasets that use Google BigQuery can only be stored in SPICE.

To create a dataset using Google BigQuery

1. Open the [Quick console](#).
2. From the start page, choose **Data**.
3. Choose **Create**, then **New Dataset**
4. Choose the **Google BigQuery** tile, and then choose **Create dataset**.
5. For **Tables**, do one of the following:

- Choose the table that you want to use.
 - Choose **Use custom SQL** to use your own personal SQL statement. For more information about using custom SQL in Quick Sight, see [Using SQL to customize data](#).
6. Choose **Edit/Preview**.
 7. (Optional) In the **Data prep** page that opens, you can add customizations to your data with calculated fields, filters, and joins.
 8. When you are finished making changes, choose **Save** to save and close the dataset.

Creating a dataset using a Google Sheets data source

Google Sheets is a web-based spreadsheet application that enables users to create, edit, and collaborate on data in real time. With its comprehensive set of functions and formulas, it serves as a powerful data source for business intelligence and analytics. Users can organize, analyze, and share insights efficiently, while its seamless collaboration features make it an ideal platform for teams working on data-driven projects.

Admin configuration in Amazon Quick

Amazon Quick administrators need to perform a one-time setup to enable Google Sheets as a data source. For detailed instructions and important considerations, see [the blog](#).

Creating a dataset using a Google Sheets data source

Use the following procedure to create a dataset using a Google Sheets data source.

To create a dataset using a Google Sheets data source

1. From the Quick start page, choose **Datasets**.
2. On the **Datasets** page, choose **New Dataset**.
3. Choose **Google Sheets**.
4. Enter a name for the data source, and then choose **Connect**.
5. When redirected to Google's sign-in page, do the following:
 - a. Enter your Google account credentials, and then choose **Next**.
 - b. Review the permissions to authorize your AWS account to connect with Google Sheets, and then choose **Continue**.

6. In the **Choose your table** menu, locate your data. The menu displays all folders, subfolders, sheets, and tabs from your Google account. To display the tabs, select a sheet from the displayed list.
7. Select the tab you want to work with.
8. Choose **Edit/Preview data** to navigate to the Data preparation page. Choose **Add data** to include any additional tabs.
9. Configure the join, and then select **Publish & visualize** to analyze your Google Sheets data with Quick Sight.

Note

- This connector supports only SPICE functionality.
- If your OAuth token expires (visible in the ingestion error report or when creating a new dataset), reauthorize by choosing **Edit** on the data source and updating it.

Creating a dataset using an Apache Impala data source

Apache Impala is a high-performance massively parallel processing (MPP) SQL query engine designed to run natively on Apache Hadoop. Use the procedure below to establish a secure connection between Quick Sight and Apache Impala.

All traffic between Quick Sight and Apache Impala is encrypted using SSL. Quick Sight supports standard username and password authentication for Impala connections.

To establish a connection, you'll need to configure SSL settings in your Impala instance, prepare your authentication credentials, set up the connection in Quick Sight using your Impala server details, and validate the connection to ensure secure data access.

To create a dataset using an Apache Impala data source

1. On the Quick start page, choose **Data**.
2. On the **Data** page, choose **Create**.
3. Choose **Data source**.
4. Choose **Impala**, then choose **Next**.
5. Enter a name for the data source.

6. For public connections:
 - a. Enter connection details for **Database server**, **HTTP Path**, **Port**, **Username**, and **Password**.
 - b. Once the validation is successful, choose **Create data source**.
7. For private connections:
 - Coordinate with your administrator to set up a VPC connection before entering connection details.

You or your administrator can [configure the VPC connection in Quick](#). SSL is enabled by default to ensure secure data transmission. If you encounter connection validation errors, please verify your connection and VPC details.

If issues persist, consult your administrator to confirm that your Certificate Authority is included in Quick Sight's [approved list of certificates](#).
8. In the **Choose your table** menu, you can either:
 - a. Choose a specific schema or table, then choose **Select**.
 - b. Choose **Use custom SQL** to write your own SQL query.
9. After completing your selection, you will be redirected to the data preparation page. Make any adjustments to your data, then choose **Publish & visualize** to analyze your Impala data in Quick Sight.

 Note

This connector supports:

- Username and password authentication
- Public and private connections
- Table discovery and custom SQL queries
- Full data refresh during ingestion
- SPICE storage only

Creating a dataset using a Microsoft Excel file

To create a dataset using a Microsoft Excel file data source, upload an .xlsx file from a local or networked drive. The data is imported into [SPICE](#).

For more information about creating new Amazon S3 datasets using Amazon S3 data sources, see [Creating a dataset using an existing Amazon S3 data source](#) or [Creating a dataset using Amazon S3 files](#).

To create a dataset based on an excel file

1. Check [Data source quotas](#) to make sure that your target file doesn't exceed data source quotas.
2. On the Quick start page, choose **Data**.
3. On the **Data** page, choose **Create** then **New dataset**.
4. Choose **Upload a file**.
5. In the **Open** dialog box, choose a file, and then choose **Open**.

A file must be 1 GB or less to be uploaded to Quick Sight.

6. If the Excel file contains multiple sheets, choose the sheet to import. You can change this later by preparing the data.

7.

Note

On the following screens, you have multiple chances to prepare the data. Each of these takes you to the **Prepare Data** screen. This screen is the same one where you can access after the data import is complete. It enables you to change the upload settings even after the upload is complete.

Choose **Select** to confirm your settings. Or you can choose **Edit/Preview data** to prepare the data immediately.

A preview of the data appears on the next screen. You can't make changes directly to the data preview.

8. If the data headings and content don't look correct, choose **Edit settings and prepare data** to correct the file upload settings.

Otherwise, choose **Next**.

9. On the **Data Source Details** screen, you can choose **Edit/Preview data**. You can specify a dataset name in the **Prepare Data** screen.

If you don't need to prepare the data, you can choose to create an analysis using the data as-is. Choose **Visualize**. Doing this names the dataset the same as the source file, and takes you to the **Analysis** screen. To learn more about data preparation and excel upload settings, see [Preparing data in Amazon Quick Sight](#).

Note

If at anytime you want to make changes to the file, such as adding a new field, you must make the change in Microsoft Excel and create a new dataset using the updated version in Quick Sight. For more information about possible implications of changing datasets, see [Things to consider when editing datasets](#).

Creating a data source using Presto

Presto (or PrestoDB) is an open-source, distributed SQL query engine, designed for fast analytic queries against data of any size. It supports both nonrelational and relational data sources. Supported nonrelational data sources include the Hadoop Distributed File System (HDFS), Amazon S3, Cassandra, MongoDB, and HBase. Supported relational data sources include MySQL, PostgreSQL, Amazon Redshift, Microsoft SQL Server, and Teradata.

For more information about Presto, see the following:

- [Introduction to presto](#), a description of Presto on the AWS website.
- [Creating a presto cluster with Amazon elastic MapReduce \(EMR\)](#) in the *Amazon EMR Release Guide*.
- For general information on Presto, see the [Presto documentation](#).

The results of the queries that you run through the Presto query engine can be turned into Quick Sight datasets. Presto processes the analytic queries on the backend databases. Then it returns results to the Quick Sight client. You can directly query your data through Presto, or you can import the results of your query into SPICE.

Before you use Quick Sight as a Presto client to run queries, make sure that you configure data source profiles. You need a data source profile in Quick Sight for each Presto data source that you want to access. Use the following procedure to create a connection to Presto.

To create a new connection to a presto data source from Amazon Quick Sight (console)

1. On the Amazon Quick Sight start page, choose **Data** at left.
2. Choose **Create** then **New dataset**.
3. Choose the **Presto** tile.

Note

In most browsers, you can use Ctrl-F or Cmd-F to open a search box and enter **presto** to locate it.

4. Add the settings for the new data source:
 - **Data source name** – Enter a descriptive name for your data source connection. This name appears in the **Existing data sources** section at the bottom of the **Data sets** screen.
 - **Connection type** – Choose the connection type that you need to use to connect to Presto.

To connect through the public network, choose **Public network**.

If you use a public network, your Presto server must be secured and authenticated using Lightweight Directory Access Protocol (LDAP). For information on configuring Presto to use LDAP, see [LDAP authentication](#) in the Presto documentation.

To connect through a virtual private connection, choose the appropriate VPC name from the **VPC connections** list.

If your Presto server allows unauthenticated access, AWS requires that you connect to it securely by using a private VPC connection. For information on configuring a new VPC, see [Configuring VPC connections in Amazon Quick Sight](#).

- **Database server** – The name of the database server.
- **Port** – The port that the server using to accept incoming connections from Amazon Quick Sight
- **Catalog** – The name of the catalog that you want to use.

- **Authentication required** – (Optional) This option only appears if you choose a VPC connection type. If the Presto data source that you're connecting to doesn't require authentication, choose **No**. Otherwise, keep the default setting (**Yes**).
- **Username** – Enter a user name to use to connect to Presto. Quick Sight applies the same user name and password to all connections that use this data source profile. If you want to monitor Quick Sight separately from other accounts, create a Presto account for each Quick Sight data source profile.

The Presto account that you use needs be able to access to the database and run SELECT statements on at least one table.

- **Password** – The password to use with the Presto user name. Amazon Quick Sight encrypts all credentials that you use in data source profile. For more information, see [Data encryption in Amazon Quick](#).
 - **Enable SSL** – SSL is enabled by default.
5. Choose **Validate connection** to test your settings.
 6. After you validate your settings, choose **Create data source** to complete the connection.

Using Snowflake with Amazon Quick Sight

Snowflake is an AI data cloud platform that provides data solutions from data warehousing and collaboration to data science and generative AI. Snowflake is an [AWS Partner](#) with multiple AWS accreditations that include AWS ISV Competencies in Generative AI, Machine Learning, Data and Analytics, and Retail.

Amazon Quick Sight offers multiple ways to connect to Snowflake: with your Snowflake login credentials, with OAuth client credentials, or with three-legged OAuth (3LO) using end-user identities. Use the following sections to learn about these methods of connection.

Topics

- [Creating an Quick Sight data source connection to Snowflake with login credentials](#)
- [Creating an Quick Sight data source connection to Snowflake with OAuth client credentials](#)
- [Creating an Quick Sight data source connection to Snowflake with three-legged OAuth \(3LO\)](#)

Creating an Quick Sight data source connection to Snowflake with login credentials

Use this section to learn how to create a connection between Quick Sight and Snowflake with your Snowflake login credentials. All traffic between Quick Sight and Snowflake is enabled by SSL.

To create a connection between Quick Sight and Snowflake

1. Open the [Quick console](#).
2. From the left navigation pane, choose **Data**, then choose **Create**, then choose **New Dataset**.
3. Choose the **Snowflake** data source card.
4. In the pop up that appears, enter the following information:
 - a. For **Data source name**, enter a descriptive name for your Snowflake data source connection. Because you can create many datasets from a connection to Snowflake, it's best to keep the name simple.
 - b. For **Connection type**, choose the type of network that you're using. Choose **Public network** if your data is shared publicly. Choose **VPC** if your data is located inside a VPC. To configure a VPC connection in Quick Sight, see [Managing VPC connection in Amazon Quick](#).
 - c. For **Database server** enter the hostname specified in your Snowflake connection details.
5. For **Database name and Warehouse**, enter the respective Snowflake database and warehouse that you want to connect.
6. For **Username** and **Password**, enter your Snowflake credentials.

After you have successfully created a data source connection between your Quick Sight and Snowflake accounts, you can begin [Creating datasets](#) that contain Snowflake data.

Creating an Quick Sight data source connection to Snowflake with OAuth client credentials

You can use OAuth client credentials to connect your Quick Sight account with Snowflake through the [Quick Sight APIs](#). *OAuth* is a standard authorization protocol that is often utilized for applications that have advanced security requirements. When you connect to Snowflake with OAuth client credentials, you can create datasets that contain Snowflake data with the Quick Sight APIs and in the Quick Sight UI. For more information about configuring OAuth in Snowflake, see [Snowflake OAuth overview](#).

Quick Sight supports the `client_credentials` OAuth grant type. OAuth client credentials is used to obtain an access token for machine-to-machine communication. This method is suitable

for scenarios where a client needs to access resources that are hosted on a server without the involvement of a user.

In the client credentials flow of OAuth 2.0, there are several client authentication mechanisms that can be used to authenticate the client application with the authorization server. Quick Sight supports client credentials based OAuth for Snowflake for the following two mechanisms:

- **Token (Client secrets-based OAuth):** The secret-based client authentication mechanism is used with the client credentials to grant flow in order to authenticate with authorization server. This authentication scheme requires the `client_id` and `client_secret` of the OAuth client app to be stored in Secrets Manager.
- **X509 (Client private key JWT-based OAuth):** The X509 certificate key-based solution provides an additional security layer to the OAuth mechanism with client certificates that are used to authenticate instead of client secrets. This method is primarily used by private clients who use this method to authenticate with the authorization server with strong trust between the two services.

Quick Sight has validated OAuth connections with the following Identity providers:

- OKTA
- PingFederate

Storing OAuth credentials in Secrets Manager

OAuth client credentials are meant for machine-to-machine use cases and are not designed to be interactive. To create a datasource connection between Quick Sight and Snowflake, create a new secret in Secrets Manager that contains your credentials for the OAuth client app. The secret ARN that is created with the new secret can be used to create datasets that contain Snowflake data in Quick Sight. For more information about using Secrets Manager keys in Quick Sight, see [Using AWS Secrets Manager secrets instead of database credentials in Quick](#).

The credentials that you need to store in Secrets Manager are determined by the OAuth mechanism that you use. The following key/value pairs are required for X509-based OAuth secrets:

- `username`: The Snowflake account username to be used when connecting to Snowflake
- `client_id`: The OAuth client ID
- `client_private_key`: The OAuth client private key

- `client_public_key`: The OAuth client certificate public key and its encrypted algorithm (for example, `{"alg": "RS256", "kid", "cert_kid"}`)

The following key/value pairs are required for token-based OAuth secrets:

- `username`: The Snowflake account username to be used when connecting to Snowflake
- `client_id`: The OAuth client ID
- `client_secret`: the OAuth client secret

Storing token-based OAuth credentials is optional

For token-based OAuth, storing these credentials in Secrets Manager is optional. Instead of creating a secret and passing its ARN, you can pass the `username`, `client_id`, and `client_secret` as inline credentials in the `CreateDataSource` request.

Creating a Snowflake OAuth connection with the Quick Sight APIs

After you create a secret in Secrets Manager that contains your Snowflake OAuth credentials and have connected your Quick account to Secrets Manager, you can establish a data source connection between Quick Sight and Snowflake with the Quick Sight APIs and SDK. The following example creates a Snowflake data source connection using token OAuth client credentials.

```
{
  "AwsAccountId": "AWSACCOUNTID",
  "DataSourceId": "UNIQUEDATASOURCEID",
  "Name": "NAME",
  "Type": "SNOWFLAKE",
  "DataSourceParameters": {
    "SnowflakeParameters": {
      "Host": "HOSTNAME",
      "Database": "DATABASENAME",
      "Warehouse": "WAREHOUSENAME",
      "AuthenticationType": "TOKEN",
      "DatabaseAccessControlRole": "snowflake-db-access-role-name",
      "OAuthParameters": {
        "TokenProviderUrl": "oauth-access-token-endpoint",
        "OAuthScope": "oauth-scope",
```

```

        "IdentityProviderResourceUri" : "resource-uri",
        "IdentityProviderVpcConnectionProperties" : {
            "VpcConnectionArn": "IdP-VPC-connection-ARN"
        }
    },
    "VpcConnectionProperties": {
        "VpcConnectionArn": "VPC-connection-ARN-for-Snowflake"
    }
    "Credentials": {
        "SecretArn": "oauth-client-secret-ARN"
    }
}

```

For more information about the CreateDataSource API operation, see [CreateDataSource](#).

Once the connection between Quick Sight and Snowflake is established and a data source is created with the Quick Sight APIs or SDK, the new data source is displayed in Quick Sight. Quick Sight authors can use this data source to create datasets that contain Snowflake data. Tables are displayed based on the role used in the DatabaseAccessControlRole parameter that is passed in a CreateDataSource API call. If this parameter is not defined when the data source connection is created, the default Snowflake role is used.

After you have successfully created a data source connection between your Quick Sight and Snowflake accounts, you can begin [Creating datasets](#) that contain Snowflake data.

Creating an Quick Sight data source connection to Snowflake with three-legged OAuth (3LO)

Quick Sight supports three-legged OAuth (3LO) with end-user identities for Snowflake.

Administrator setup: Client application configuration

The Quick Sight account administrator must configure a client application before authors can create 3LO data sources. Complete the following steps:

1. In the Quick Sight console, choose **Manage Account**. Under **Security**, choose **Manage OAuth Client Applications**.
2. Choose **Add OAuth client application** and provide the following information:
 - a. Config ID
 - b. OAuth client application name

- c. Authentication type
 - d. Client ID
 - e. Client secret
 - f. Token endpoint URL
 - g. Authorization endpoint URL
 - h. OAuth scopes
 - i. VPC connection ARN (optional)
 - j. Data source type (select **SNOWFLAKE**)
3. Choose **Add** to save the client application.

Author setup: Creating a 3LO data source

After the administrator configures the client application, authors can create a Snowflake data source with 3LO authentication.

1. Open the [Quick console](#).
2. Choose **Datasets**, then choose **New dataset** to open the **Create Data Source** page.
3. Choose **Snowflake**.
4. Choose **3LO** as the authentication method.
5. Enter the required connection information.
6. Choose **Create data source**.

Using Starburst with Amazon Quick Sight

Starburst is a full-featured data lake analytics service built on top of a massively parallel processing (MPP) query engine, Trino. Use this section to learn how to connect from Amazon Quick Sight to Starburst. All traffic between Quick Sight and Starburst is enabled by SSL. If you're connecting to Starburst Galaxy, you can get the necessary connection details by logging in to your Starburst Galaxy account, then choose **Partner Connect** and then **Quick Sight**. You should be able to see information, such as hostname and port. Amazon Quick Sight supports basic username and password authentication to Starburst.

Quick Sight offers two ways to connect to Starburst: with your Starburst login credentials or with OAuth client credentials. Use the following sections to learn about both methods of connection.

Topics

- [Creating an Quick Sight data source connection to Starburst with login credentials](#)
- [Creating an Quick Sight data source connection to Starburst with OAuth client credentials](#)

Creating an Quick Sight data source connection to Starburst with login credentials

1. Begin by creating a new dataset. From the left navigation pane, choose **Data**, then choose **Create**, then choose **New Dataset**.
2. Choose the **Starburst** data source card.
3. Select the Starburst product type. Choose **Starburst Enterprise** for on-prem Starburst instances. Choose **Starburst Galaxy** for managed instances.
4. For **Data source name**, enter a descriptive name for your Starburst data source connection. Because you can create many datasets from a connection to Starburst, it's best to keep the name simple.
5. For **Connection type**, select the type of network you're using. Choose **Public network** if your data is shared publicly. Choose **VPC** if your data is inside a VPC. To configure a VPC connection in Amazon Quick Sight, see [Configuring the VPC connection in Amazon Quick Sight](#). This connection type is not available for Starburst Galaxy.
6. For **Database server** enter the hostname specified in your Starburst connection details.
7. For **Catalog**, enter the catalog specified in your Starburst connection details.
8. For **Port**, enter the port specified in your Starburst connection details. Defaults to 443 for Starburst Galaxy.
9. For **Username** and **Password**, enter your Starburst connection credentials.
10. To verify the connection is working, choose **Validate connection**.
11. To finish and create the data source, choose **Create data source**.

Note

Connectivity between Amazon Quick Sight and Starburst was validated using Starburst version 420.

After you have successfully created a data source connection between your Quick Sight and Starburst accounts, you can begin [Creating datasets](#) that contain Starburst data.

Creating an Quick Sight data source connection to Starburst with OAuth client credentials

You can use OAuth client credentials to connect your Quick Sight account with Starburst through the [Quick Sight APIs](#). *OAuth* is a standard authorization protocol that is often utilized for applications that have advanced security requirements. When you connect to Starburst with OAuth client credentials, you can create datasets that contain Starburst data with the Quick Sight APIs and in the Quick Sight UI. For more information about configuring OAuth in Starburst, see [OAuth 2.0 authentication](#).

Quick Sight supports the `client_credentials` OAuth grant type. OAuth client credentials is used to obtain an access token for machine-to-machine communication. This method is suitable for scenarios where a client needs to access resources that are hosted on a server without the involvement of a user.

In the client credentials flow of OAuth 2.0, there are several client authentication mechanisms that can be used to authenticate the client application with the authorization server. Quick Sight supports client credentials based OAuth for Starburst for the following two mechanisms:

- **Token (Client secrets-based OAuth):** The secret-based client authentication mechanism is used with the client credentials to grant flow in order to authenticate with authorization server. This authentication scheme requires the `client_id` and `client_secret` of the OAuth client app to be stored in Secrets Manager.
- **X509 (Client private key JWT-based OAuth):** The X509 certificate key-based solution provides an additional security layer to the OAuth mechanism with client certificates that are used to authenticate instead of client secrets. This method is primarily used by private clients who use this method to authenticate with the authorization server with strong trust between the two services.

Quick Sight has validated OAuth connections with the following Identity providers:

- OKTA
- PingFederate

Storing OAuth credentials in Secrets Manager

OAuth client credentials are meant for machine-to-machine use cases and are not designed to be interactive. To create a datasource connection between Quick Sight and Starburst, create a new secret in Secrets Manager that contains your credentials for the OAuth client app. The secret ARN

that is created with the new secret can be used to create datasets that contain Starburst data in Quick Sight. For more information about using Secrets Manager keys in Quick Sight, see [Using AWS Secrets Manager secrets instead of database credentials in Quick](#).

The credentials that you need to store in Secrets Manager are determined by the OAuth mechanism that you use. The following key/value pairs are required for X509-based OAuth secrets:

- `username`: The Starburst account username to be used when connecting to Starburst
- `client_id`: The OAuth client ID
- `client_private_key`: The OAuth client private key
- `client_public_key`: The OAuth client certificate public key and its encrypted algorithm (for example, `{"alg": "RS256", "kid", "cert_kid"}`)

The following key/value pairs are required for token-based OAuth secrets:

- `username`: The Starburst account username to be used when connecting to Starburst
- `client_id`: The OAuth client ID
- `client_secret`: the OAuth client secret

Creating a Starburst OAuth connection with the Quick Sight APIs

After you create a secret in Secrets Manager that contains your Starburst OAuth credentials and have connected your Quick account to Secrets Manager, you can establish a data source connection between Quick Sight and Starburst with the Quick Sight APIs and SDK. The following example creates a Starburst data source connection using token OAuth client credentials.

```
{
  "AwsAccountId": "AWSACCOUNTID",
  "DataSourceId": "DATASOURCEID",
  "Name": "NAME",
  "Type": "STARBURST",
  "DataSourceParameters": {
    "StarburstParameters": {
      "Host": "STARBURST_HOST_NAME",
      "Port": "STARBURST_PORT",
      "Catalog": "STARBURST_CATALOG",
      "ProductType": "STARBURST_PRODUCT_TYPE",
      "AuthenticationType": "TOKEN",
      "DatabaseAccessControlRole": "starburst-db-access-role-name",
```



```

    "OAuthParameters": {
      "TokenProviderUrl": "oauth-access-token-endpoint",
      "OAuthScope": "oauth-scope",
      "IdentityProviderResourceUri" : "resource-uri",
      "IdentityProviderVpcConnectionProperties" : {
        "VpcConnectionArn": "IdP-VPC-connection-ARN"
      }
    },
    "VpcConnectionProperties": {
      "VpcConnectionArn": "VPC-connection-ARN-for-Starburst"
    },
    "Credentials": {
      "SecretArn": "oauth-client-secret-ARN"
    }
  }
}

```

For more information about the CreateDataSource API operation, see [CreateDataSource](#).

Once the connection between Quick Sight and Starburst is established and a data source is created with the Quick Sight APIs or SDK, the new data source is displayed in Quick Sight. Quick Sight authors can use this data source to create datasets that contain Starburst data. Tables are displayed based on the role used in the DatabaseAccessControlRole parameter that is passed in a CreateDataSource API call. If this parameter is not defined when the data source connection is created, the default Starburst role is used.

After you have successfully created a data source connection between your Quick Sight and Starburst accounts, you can begin [Creating datasets](#) that contain Starburst data.

Creating a data source and data set from SaaS sources

To analyze and report on data from software as a service (SaaS) applications, you can use SaaS connectors to access your data directly from Quick Sight. The SaaS connectors simplify accessing third-party application sources using OAuth, without any need to export the data to an intermediate data store.

You can use either a cloud-based or server-based instance of a SaaS application. To connect to an SaaS application that is running on your corporate network, make sure that Quick Sight can access the application's Domain Name System (DNS) name over the network. If Quick Sight can't access the SaaS application, it generates an unknown host error.

Here are examples of some ways that you can use SaaS data:

- Engineering teams who use Jira to track issues and bugs can report on developer efficiency and bug burndown.
- Marketing organizations can integrate Quick Sight with Adobe Analytics to build consolidated dashboards to visualize their online and web marketing data.

Use the following procedure to create a data source and dataset by connecting to sources available through Software as a Service (SaaS). In this procedure, we use a connection to GitHub as an example. Other SaaS data sources follow the same process, although the screens—especially the SaaS screens—might look different.

To create a data source and dataset by connecting to sources through SaaS

1. On the Quick start page, choose **Data**.
2. On the **Data** page, choose **Create** then choose **New dataset**.
3. Choose the icon that represents the SaaS source that you want to use. For example, you might choose Adobe Analytics or GitHub.

For sources using OAuth, the connector takes you to the SaaS site to authorize the connection before you can create the data source.

4. Choose a name for the data source, and enter that. If there are more screen prompts, enter the appropriate information. Then choose **Create data source**.
5. If you are prompted to do so, enter your credentials on the SaaS login page.
6. When prompted, authorize the connection between your SaaS data source and Quick Sight.

The following example shows the authorization for Quick Sight to access the GitHub account for the Quick Sight documentation.

Note

Quick Sight documentation is now available on GitHub. If you want to make changes to this user guide, you can use GitHub to edit it directly.

(Optional) If your SaaS account is part of an organizational account, you might be asked to request organization access as part of authorizing Quick Sight. If you want to do this, follow the prompts on your SaaS screen, then choose to authorize Quick Sight.

7. After authorization is complete, choose a table or object to connect to. Then choose **Select**.
8. On the **Finish data set creation** screen, choose one of these options:
 - To save the data source and dataset, choose **Edit/Preview data**. Then choose **Save** from the top menu bar.
 - To create a dataset and an analysis using the data as-is, choose **Visualize**. This option automatically saves the data source and the dataset.

You can also choose **Edit/Preview data** to prepare the data before creating an analysis. This opens the data preparation screen. For more information about data preparation, see [Preparing dataset examples](#).

The following constraints apply:

- The SaaS source must support REST API operations for Quick Sight to connect to it.
- If you are connecting to Jira, the URL must be public address.
- If you don't have enough [SPICE](#) capacity, choose **Edit/Preview data**. In the data preparation screen, you can remove fields from the dataset to decrease its size or apply a filter that reduces the number of rows returned. For more information about data preparation, see [Preparing dataset examples](#).

Creating a dataset from Salesforce

Use the following procedure to create a dataset by connecting to Salesforce and selecting a report or object to provide data.

To create a dataset using Salesforce from a report or object

1. Check [Data source quotas](#) to make sure that your target report or object doesn't exceed data source quotas.
2. On the Quick start page, choose **Data**.
3. On the **Data** page, choose **Create** then **New dataset**.
4. Choose the **Salesforce** icon.
5. Enter a name for the data source and then choose **Create data source**.
6. On the Salesforce login page, enter your Salesforce credentials.

7. For **Data elements: contain your data**, choose **Select** and then choose either **REPORT** or **OBJECT**.

 Note

Joined reports aren't supported as Quick Sight data sources.

8. Choose one of the following options:
 - To prepare the data before creating an analysis, choose **Edit/Preview data** to open data preparation. For more information about data preparation, see [Preparing dataset examples](#).
 - Otherwise, choose a report or object and then choose **Select**.
9. Choose one of the following options:
 - To create a dataset and an analysis using the data as-is, choose **Visualize**.

 Note

If you don't have enough [SPICE](#) capacity, choose **Edit/Preview data**. In data preparation, you can remove fields from the dataset to decrease its size or apply a filter that reduces the number of rows returned. For more information about data preparation, see [Preparing dataset examples](#).

- To prepare the data before creating an analysis, choose **Edit/Preview data** to open data preparation for the selected report or object. For more information about data preparation, see [Preparing dataset examples](#).

 Note

The Salesforce connector is not supported in embedded console deployments where users authenticate through namespace isolation. The OAuth authentication flow requires direct Amazon Quick Sight console access to complete the sign-in process.

Using Trino with Amazon Quick Sight

Trino is a massively parallel processing (MPP) query engine built to quickly query data lakes with petabytes of data. Use this section to learn how to connect from Amazon Quick Sight to Trino. All

traffic between Amazon Quick Sight and Trino is enabled by SSL. Amazon Quick Sight supports basic username and password authentication to Trino.

Creating a data source connection for Trino

1. Begin by creating a new dataset. From the left navigation pane, choose **Data**. Choose **Create** then **New Dataset**.
2. Choose the **Trino** data source card.
3. For **Data source name**, enter a descriptive name for your Trino data source connection. Because you can create many datasets from a connection to Trino, it's best to keep the name simple.
4. For **Connection type**, select the type of network you're using. Choose **Public network** if your data is shared publicly. Choose **VPC** if your data is inside a VPC. To configure a VPC connection in Amazon Quick Sight, see [Configuring the VPC connection in Amazon Quick Sight](#).
5. For **Database server**, enter the hostname specified in your Trino connection details.
6. For **Catalog**, enter the catalog specified in your Trino connection details.
7. For **Port**, enter the port specified in your Trino connection details.
8. For **Username** and **Password**, enter your Trino connection credentials.
9. To verify the connection is working, choose **Validate connection**.
10. To finish and create the data source, choose **Create data source**.

Adding a new Amazon Quick Sight dataset for Trino

After you go through the [data source creation process](#) for Trino, you can create Trino datasets to use for analysis. You can create new datasets from a new or an existing Trino data source. When you are creating a new data source, Amazon Quick Sight immediately takes you to creating a dataset, which is step 3 below. If you're using an existing data source to create a new dataset, start from step 1 below.

To create a dataset using a Trino data source, see the following steps.

1. From the start page, choose **Data**. Choose **Create** then **New dataset**.
2. Choose the Trino data source you created.
3. Choose **Create data set**.
4. To specify the table you want to connect to, choose a schema. If you don't want to choose a schema, you can also use your own SQL statement.

5. To specify the table you want to connect to, first select the **Schema** you want to use. For **Tables**, choose the table that you want to use. If you prefer to use your own SQL statement, select **Use custom SQL**.
6. Choose **Edit/Preview**.
7. (Optional) To add more data, use the following steps:
8. Choose **Add data** in the top right.
9. To connect to different data, choose **Switch data source**, and choose a different dataset.
10. Follow the prompts to finish adding data.
11. After adding new data to the same dataset, choose **Configure this join** (the two red dots). Set up a join for each additional table.
12. If you want to add calculated fields, choose **Add calculated field**.
13. Clear the check box for any fields that you want to omit.
14. Update any data types that you want to change.
15. When you are done, choose **Save** to save and close the dataset.

Note

Connectivity between Quick Sight and Trino was validated using Trino version 410.

Creating a dataset using a local text file

To create a dataset using a local text file data source, identify the location of the file, and then upload it. The file data is automatically imported into [SPICE](#) as part of creating a dataset.

To create a dataset based on a local text file

1. Check [Data source quotas](#) to make sure that your target file doesn't exceed data source quotas.

Supported file types include .csv, .tsv, .json, .clf, or .elf files.

2. On the Quick start page, choose **Data**.
3. Choose **Create** then **New dataset**.
4. Choose **Upload a file**.
5. In the **Open** dialog box, browse to a file, select it, and then choose **Open**.

A file must be 1 GB or less to be uploaded to Quick Sight.

6. To prepare the data before creating the dataset, choose **Edit/Preview data**. Otherwise, choose **Visualize** to create an analysis using the data as-is.

If you choose the former, you can specify a dataset name as part of preparing the data. If you choose the latter, a dataset with the same name as the source file is created. To learn more about data preparation, see [Preparing data in Amazon Quick Sight](#).

Using Amazon Timestream data with Amazon Quick Sight

Following, you can find how to connect to your Amazon Timestream data using Amazon Quick Sight. For a brief overview, see the [Getting started with Amazon Timestream and Amazon QuickSight](#) video tutorial on YouTube.

Creating a new Amazon Quick Sight data source connection for a Timestream database

Following, you can find how to connect to Amazon Timestream from Amazon Quick Sight.

Before you can proceed, Amazon Quick Sight needs to be authorized to connect to Amazon Timestream. If connections aren't enabled, you get an error when you try to connect. A Quick Sight administrator can authorize connections to AWS resources. To authorize, open the menu by clicking on your profile icon at top right. Choose **Manage QuickSight, Security & permissions, Add or remove**. Then enable the check box for Amazon Timestream, then choose **Update** to confirm. For more information, see [Configuring Amazon Quick Sight access to AWS data sources](#).

To connect to Amazon Timestream

1. Begin by creating a new dataset. Choose **Data** from the navigation pane at left.
2. Choose **Create** then **New Dataset**.
3. Choose the Timestream data source card.
4. For **Data source name**, enter a descriptive name for your Timestream data source connection, for example US Timestream Data. Because you can create many datasets from a connection to Timestream, it's best to keep the name simple.
5. Choose **Validate connection** to check that you can successfully connect to Timestream.
6. Choose **Create data source** to proceed.
7. For **Database**, choose **Select** to view the list of available options.
8. Choose the one you want to use, then choose **Select** to continue.

9. Do one of the following:

- To import your data into Quick Sight's in-memory engine (called SPICE), choose **Import to SPICE for quicker analytics**.
- To allow Quick Sight to run a query against your data each time you refresh the dataset or use the analysis or dashboard, choose **Directly query your data**.

If you want to enable autorefresh on a published dashboard that uses Timestream data, the Timestream dataset needs to use a direct query.

10. Choose **Edit/Preview** and then **Save** to save your dataset and close it.

11. Repeat these steps for the number of concurrent direct connections to Timestream that you want to open in a dataset. For example, let's say you want to use four tables in a Quick Sight dataset. Currently, Quick Sight datasets connect to only one table at a time from a Timestream data source. To use four tables in the same dataset, you need to add four data source connections in Quick Sight.

Managing permissions for Timestream data

The following procedure describes how to view, add, and revoke permissions to allow access to the same Timestream data source. The people that you add need to be active users in Quick Sight before you can add them.

To edit permissions on a dataset

1. Choose **Data** at left, then scroll down to find the dataset for your Timestream connection. An example might be **US Timestream Data**.
2. Choose the **Timestream** dataset to open it.
3. On the dataset details page that opens, choose the **Permissionstab**.

A list of current permissions appears.

4. To add permissions, choose **Add users & groups**, then follow these steps:
 - a. Add users or groups to allow them to use the same dataset.
 - b. When you're finished adding everyone that you want to add, choose the **Permissions** that you want to apply to them.
5. (Optional) To edit permissions, you can choose **Viewer** or **Owner**.

- Choose **Viewer** to allow read access.
 - Choose **Owner** to allow that user to edit, share, or delete this Quick Sight data source.
6. (Optional) To revoke permissions, choose **Revoke access**. After you revoke someone's access, they can't create edit, share, or delete the dataset.
 7. When you are finished, choose **Close**.

Adding a new Quick Sight dataset for Timestream

After you have an existing data source connection for Timestream data, you can create Timestream datasets to use for analysis.

Currently, you can use a Timestream connection only for a single table in a dataset. To add data from multiple Timestream tables in a single dataset, create an additional Quick Sight data source connection for each table.

To create a dataset using Amazon Timestream

1. Choose **Data** at left, then scroll down to find the data source card for your Timestream connection. If you have many data sources, you can use the search bar at the top of the page to find your data source with a partial match on the name.
2. Choose the **Timestream** data source card, and then choose **Create data set**.
3. For **Database**, choose **Select** to view a list of available databases and choose the one that you want to use.
4. For **Tables**, choose the table that you want to use.
5. Choose **Edit/Preview**.
6. (Optional) To add more data, use the following steps:
 - a. Choose **Add data** at top right.
 - b. To connect to different data, choose **Switch data source**, and choose a different dataset.
 - c. Follow the UI prompts to finish adding data.
 - d. After adding new data to the same dataset, choose **Configure this join** (the two red dots). Set up a join for each additional table.
 - e. If you want to add calculated fields, choose **Add calculated field**.
 - f. To add a model from SageMaker AI, choose **Augment with SageMaker**. This option is only available in Amazon Quick Enterprise edition.

- g. Clear the check box for any fields that you want to omit.
 - h. Update any data types that you want to change.
7. When you are done, choose **Save** to save and close the dataset.

Adding Timestream data to an analysis

Following, you can find how to add an Amazon Timestream dataset to a Quick Sight analysis. Before you begin, make sure that you have an existing dataset that contains the Timestream data that you want to use.

To add Amazon Timestream data to an analysis

1. Choose **Analyses** at left.
2. Do one of the following:
 - To create a new analysis, choose **New analysis** at right.
 - To add to an existing analysis, open the analysis that you want to edit.
 - Choose the pencil icon near at top left.
 - Choose **Add data set**.
3. Choose the Timestream dataset that you want to add.

For more information, see [Working with analyses](#).

Connecting to data catalogs in Amazon Quick

Note

Data catalog integration is in Preview for Amazon Quick and is subject to change. Preview features are not typically recommended for critical or production workloads.

Amazon Quick data catalog integration connects Quick to upstream data catalogs such as AWS Glue Data Catalog and Databricks Unity Catalog. The integration provides an AI-powered agentic experience that helps data curators (authors such as analysts and business intelligence engineers) rapidly discover relevant tables, bulk-create DirectQuery datasets and topics, and automatically inherit semantic definitions from their catalogs. By reusing the metadata that already exists in your

catalog, you can go from a business question to production-ready dashboards and AI-powered Q&A in a fraction of the time, not weeks.

In Quick, a connected catalog is represented as a data source. To access a catalog data source, navigate to the **Create Data Source** page and select the catalog connector for your environment.

Benefits of data catalog integration

Data catalog integration provides the following benefits:

- **Accelerated asset curation** – Data curators (authors) use a conversational AI agent to discover tables that match their business needs, create datasets in bulk, and build multi-dataset topics in a single flow. These datasets are DirectQuery representations of your catalog assets, not copies of your data. The upstream catalog remains the source of truth. This significantly reduces the time required to deliver governed analytics to business users.
- **Semantic inheritance** – Table and column descriptions, and primary and foreign key relationships (where available), are automatically inherited from the upstream catalog. This eliminates manual recreation of business context and ensures consistency with your catalog as the single source of truth.
- **Focused context boundaries** – In enterprise environments with thousands of tables, the agentic experience helps authors define a curated subset of catalog assets relevant to their line of business. These curated datasets and topics provide a focused data boundary that improves the accuracy of AI-generated answers in Q&A and enables deterministic, reliable dashboards.
- **Always-fresh data** – All datasets created through the catalog integration use DirectQuery, so data is queried at the source with no duplication or staleness.
- **Governance inheritance** – Optionally enable identity propagation to enforce per-user data permissions defined in your upstream catalog at query time, without recreating row-level security (RLS) or column-level security (CLS) rules in Quick.
- **Combined enterprise context** – After you establish a context boundary with datasets and topics, Quick can combine catalog data with other enterprise knowledge sources such as Slack conversations, Outlook emails, and Google Drive documents. This unified context enables AI agents to deliver more comprehensive insights across structured and unstructured data.

Supported catalogs

The following table lists the supported catalogs and their capabilities in this preview.

Supported catalogs and capabilities

Catalog	Capability	What it enables	Authentication
AWS Glue Data Catalog	Agentic discovery, bulk create, semantic inheritance	AI-powered discovery and dataset creation from Glue metadata, with automatic inheritance of table and column descriptions	Service role or AWS IAM Identity Center
AWS Glue Data Catalog	Trusted identity propagation (TIP)	Per-user data permissions enforcement through AWS IAM Identity Center and Lake Formation	AWS IAM Identity Center
Databricks Unity Catalog	Agentic discovery, bulk create, semantic inheritance	AI-powered discovery and dataset creation from Unity Catalog metadata, with automatic inheritance of descriptions and relationships	Personal Access Token (PAT) or OAuth 3LO
Databricks Unity Catalog	Identity propagation (3LO)	Per-user data permissions enforcement through Databricks OAuth	OAuth 3LO

The agentic catalog experience

After you connect to a supported catalog, data curators (authors) can choose **Explore data** on the data source to launch the agentic experience. The AI agent guides you through the following steps in a conversational flow:

1. **Discover** – Describe your use case in natural language. The agent uses catalog metadata, including table and column descriptions, data quality scores, and lineage, to recommend the most relevant tables. For example: *"I am a supply chain analyst looking to enable Q&A for my logistics managers. They need to track carrier on-time delivery performance and shipping costs across distribution centers."*
2. **Create** – After you verify the recommended assets, create all relevant datasets in bulk within the same conversational flow. These datasets are DirectQuery representations of your catalog assets — the upstream catalog remains the source of truth. We recommend that you create datasets only for your intended use cases to maintain a focused context boundary.

3. **Relationships and topics** – Once datasets are created, the agent inherits relationships from the upstream catalog (where available) and can recommend additional inferred relationships. The agent then creates a multi-dataset topic that spans your selected assets.
4. **Semantic inheritance** – Table and column descriptions are automatically inherited from the upstream catalog onto the created datasets. This metadata provides the business context that AI agents in Quick need to generate accurate, grounded answers for your end users.

Important

The agentic experience is AI-powered. Authors should review all agent recommendations, including discovered tables, inferred relationships, and inherited descriptions, before proceeding.

Behavior of created datasets

Datasets created through the agentic catalog experience have the following characteristics:

- **Semantics inherited indicator** – Created datasets display a **Semantics inherited** badge to show that definitions come from the upstream catalog.
- **DirectQuery** – All datasets use DirectQuery mode by default. Data is always queried at the source with no duplication.
- **Sync semantics** – Choose **Sync semantics** on a dataset to refresh table and column definitions from the upstream catalog at any time. This ensures the catalog remains the single source of truth for metadata.
- **Flexibility** – You can edit these datasets to switch to SPICE or add transformations in Quick. However, after you do so, semantic sync is no longer available and the dataset is treated as a standard Quick dataset.

Note

In this preview, Quick inherits table and column descriptions, and primary and foreign key definitions. Support for additional semantic information such as metric views and semantic views is planned for the future.

Identity propagation for data governance

In addition to the agentic experience and semantic inheritance, Quick supports identity propagation for applicable catalogs. When enabled, identity propagation enforces per-user data permissions at query time based on the user's identity in the upstream catalog.

For example, if John has access to US-Northeast region data and Mary has access to US-West region data, when each user views the same dashboard, they see only the data they are permitted to access. Permissions are enforced by the upstream catalog, not by manual Quick configuration.

Important

Identity propagation works only for DirectQuery datasets. If a dataset is switched to SPICE or has transformations applied, identity propagation does not apply.

Important

The agentic catalog experience (natural language discovery, bulk creation of datasets and topics, semantic inheritance) works with any supported authentication type. Identity propagation is not a prerequisite for the agentic experience. It is an optional, additive capability for customers who require per-user governance enforcement.

If you do not enable identity propagation, you can manage data permissions manually by using row-level security (RLS) and column-level security (CLS) rules in Quick.

The following identity propagation methods are available:

- **AWS Glue Data Catalog** – Trusted identity propagation through AWS IAM Identity Center and Lake Formation. For setup instructions, see [Setting up trusted identity propagation for AWS Glue Data Catalog](#).
- **Databricks Unity Catalog** – OAuth 3LO (three-legged OAuth) with end-user identities. For setup instructions, see [Setting up identity enforcement with Databricks 3LO](#).

When to use data catalog integration

Data catalog integration is designed for the following scenarios:

- **Large-scale catalogs that require a focused context boundary** – In enterprise environments with thousands of tables, exposing the entire catalog to end users can introduce noise and reduce answer accuracy in AI-powered Q&A. The agentic experience enables authors to define a curated subset of assets relevant to their line of business. These assets are represented as datasets and topics in Quick, providing a focused boundary for deterministic dashboards and more accurate AI-generated answers.
- **Catalog as the authoritative source of truth** – Datasets created through the agentic experience are DirectQuery projections of upstream catalog assets. The catalog remains the single source of truth for metadata and definitions. Semantic sync keeps definitions current. Quick also supports additional transformations on these assets when needed.
- **Combined enterprise context** – By building curated representations (datasets and topics) in Quick, Quick can combine catalog data with other enterprise knowledge sources. This unified context enables AI agents to deliver comprehensive insights across structured and unstructured data.

Connecting to AWS Glue Data Catalog

Note

Data catalog integration with AWS Glue Data Catalog is in Preview for Amazon Quick and is subject to change. Preview features are not typically recommended for critical or production workloads.

AWS Glue Data Catalog is available as a data source connector in Quick. It provides direct access to Glue catalog metadata including table names, descriptions, column information, and data quality scores. After you connect, data curators (authors) can use the agentic experience to discover relevant tables, create DirectQuery datasets in bulk, and inherit semantic definitions automatically.

Setting up the Data Catalog connection

Connecting to AWS Glue Data Catalog involves two steps: connecting to the catalog for metadata access and connecting to Amazon Athena for data queries.

Prerequisites

- An administrator must enable the AWS Glue Data Catalog connector. To do this, go to **AWS resources** under the **Manage Account** page in Quick and select **Glue Data Catalog**.

Important

If an author attempts to create an AWS Glue Data Catalog data source before the administrator enables it under **AWS resources** on the **Manage Account** page, the operation fails with an error. Ensure that the administrator enables AWS Glue Data Catalog for the Quick account before authors attempt to create the data source.

- The connector uses service role permissions for Glue metadata access. Attach the following IAM policy to the Quick service role:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "glue:GetCatalog",
        "glue:GetCatalogs",
        "glue:GetDatabases",
        "glue:GetDatabase",
        "glue:GetTables",
        "glue:GetTable",
        "glue:ListDataQualityResults",
        "glue:GetDataQualityResult",
        "lakeformation:GetTemporaryGlueTableCredentials",
        "lakeformation:GetDataAccess",
        "glue:GetConnection",
        "glue:GetConnections",
        "glue:GetTags"
      ],
      "Resource": [
        "arn:aws:glue:*:account-id:catalog",
        "arn:aws:glue:*:account-id:catalog/*",
        "arn:aws:glue:*:account-id:database/*",
        "arn:aws:glue:*:account-id:table/*/*",
        "arn:aws:glue:*:account-id:dataQualityRuleset/*",
        "arn:aws:glue:*:account-id:connection/*"
      ]
    }
  ]
}
```



```
}
```

Replace *account-id* with your AWS account ID.

Note

This policy grants access to all Glue Data Catalog resources in your account. To follow the principle of least privilege, replace the wildcard (*) resource ARNs with ARNs for the specific catalogs, databases, and tables that you want to expose in Quick.

- If you are not using AWS IAM Identity Center, all data source creators have the same access to catalog metadata as the Quick service role.

Two-step connection process

Because AWS Glue Data Catalog provides metadata and Amazon Athena provides query execution, authors complete two connections:

1. **Connect to AWS Glue Data Catalog for metadata** – Navigate to **Create Data Source** and select **Glue Data Catalog**. This connection provides access to table names, table descriptions, column information, and other catalog metadata.
2. **Connect to Amazon Athena for data queries** – Select **Athena** and configure the primary workgroup. This connection enables queries against the underlying Amazon S3 and Apache Iceberg tables. The supported underlying data store is Amazon S3.

Using the agentic experience with AWS Glue Data Catalog

After you connect to AWS Glue Data Catalog, choose **Explore data** on the data source to launch the agentic experience.

The agentic experience provides the following capabilities for Data Catalog:

- **Natural language discovery** – Describe your use case and the AI agent finds relevant tables from your catalog metadata.
- **Bulk create assets** – Create multiple datasets in a single conversational flow.
- **Semantic inheritance** – Table and column descriptions are automatically inherited from Glue onto the created datasets.

Note

AWS Glue Data Catalog does not store primary and foreign key relationships. However, the agent can recommend inferred relationships to create multi-dataset topics. Review these recommendations before accepting them.

For details on the full agentic flow including discovery, creation, relationships, and semantic inheritance, see [The agentic catalog experience](#).

Setting up trusted identity propagation for AWS Glue Data Catalog

Trusted identity propagation enables AWS services to access resources based on the end user's identity context. When you configure Quick, AWS Glue Data Catalog, Amazon S3 Access Grants, and Lake Formation with AWS IAM Identity Center, the signed-in user's identity is propagated across services. Lake Formation and S3 Access Grants make authorization decisions based on permissions defined for that user or their group membership from your identity provider.

Important

Trusted identity propagation with AWS Glue Data Catalog requires AWS IAM Identity Center. It is not supported for other authentication methods.

Prerequisites**Important**

Your AWS IAM Identity Center instance, Lake Formation, AWS Glue Data Catalog, and Amazon S3 Access Grants must all be deployed in the same AWS Region.

- **Configure Quick with AWS IAM Identity Center** – Trusted identity propagation is supported only for Quick accounts integrated with AWS IAM Identity Center. To create AWS Glue Data Catalog data sources with identity propagation, you must be an AWS IAM Identity Center user (author).
- **AWS Glue Data Catalog must be managed by Lake Formation** – Lake Formation must enforce permissions on your databases and tables. If a database or table still grants Super to IAMAllowedPrincipals, Lake Formation does not enforce per-user permissions.

- **Lake Formation must be integrated with AWS IAM Identity Center** – Use the same AWS IAM Identity Center instance that Quick uses. This integration allows Lake Formation to recognize a propagated user identity.
- **Register data S3 locations with Lake Formation** – When a user queries a table, Lake Formation vends temporary, scoped S3 credentials for the data. The user's query does not need its own S3 permissions on the data bucket. Register the location for all table formats (Parquet, Apache Iceberg, S3 Tables).
- **Grant permissions to IAM Identity Center users and groups in Lake Formation** – Configure permissions and any row, column, or cell data filters to control which catalogs, databases, tables, rows, and columns each user can access.
- **Quick administrator authorizes the connection** – With trusted identity propagation, the Quick role does not need direct S3 permissions on the table data. Lake Formation vends those credentials.
- **S3 Access Grants on the Athena results bucket** – Glue queries execute through the Athena engine, which writes results to the query results bucket. That bucket is governed by S3 Access Grants (separate from Lake Formation). Each user or group may need an S3 Access Grant (READ/ WRITE) on the results bucket.

Step 1: Enable the Lake Formation – AWS IAM Identity Center integration

Lake Formation must be connected to your AWS IAM Identity Center instance so it can recognize identities propagated from Quick. This is a one-time, account-level and Region-level step. If you already use Athena trusted identity propagation, this step is already complete. You must be a Lake Formation data lake administrator.

Console: In the Lake Formation console, choose **Administration, IAM Identity Center integration**, and then **Create**. Select your AWS IAM Identity Center instance.

AWS CLI:

```
aws lakeformation create-lake-formation-identity-center-configuration \
  --catalog-id 111122223333 \
  --instance-arn arn:aws:sso:::instance/<idc-instance-id> \
  --region us-west-2
```

This command creates the Lake Formation AWS IAM Identity Center application. Note the `ApplicationArn` from the response. You use it as the authorized target when you configure identity propagation.

To retrieve the application ARN later, run:

```
aws lakeformation describe-lake-formation-identity-center-configuration \
  --catalog-id 111122223333 \
  --region us-west-2
```

The response includes an `ApplicationArn` in the following format:

```
arn:aws:sso::111122223333:application/<idc-instance-id>/<lf-application-id>
```

Step 2: Configure the IAM role with required permissions

Your Quick account must use an IAM role with the required permissions. If your account already uses a custom IAM role, modify it. Otherwise, create a new role.

Required trust policy:

The role's trust policy must allow the Quick service principal to assume it and attach the user's identity context to the session:

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow",
    "Principal": { "Service": "quicksight.amazonaws.com" },
    "Action": ["sts:AssumeRole", "sts:TagSession", "sts:SetContext"]
  }]
}
```

Note

`sts:SetContext` is required so that the user's AWS IAM Identity Center identity context can be attached to the assumed-role session.

Required permissions:

The role needs permissions to read AWS Glue Data Catalog metadata and obtain Lake Formation vended data access. The role does not need direct `s3:*` permissions on the table data because Lake Formation vends scoped, temporary S3 credentials.

```
"Action": [  
    "glue:GetCatalog", "glue:GetCatalogs",  
    "glue:GetDatabase", "glue:GetDatabases",  
    "glue:GetTable", "glue:GetTables",  
    "glue:GetConnection", "glue:GetConnections",  
    "glue:GetTags", "glue:ListDataQualityResults",  
    "glue:GetDataQualityResult",  
    "lakeformation:GetTemporaryGlueTableCredentials",  
    "lakeformation:GetDataAccess"  
]
```

Step 3: Configure the Quick account to use the IAM role

This step must be completed by the Quick administrator.

1. In the Quick console, choose your profile icon and then choose **Manage Quick**.
2. Choose **Security & permissions**.
3. Under **Quick access to AWS services**, choose **Manage**.
4. Find or add the **Glue Data Catalog** resource.
5. Select the IAM role that you configured in Step 2.
6. Under **Amazon S3**, allow access to the Athena query results bucket.
7. Choose **Update**.

Step 4: Update the identity propagation configuration

Run the following AWS CLI command to enable identity propagation for Data Catalog:

```
aws quicksight update-identity-propagation-config \  
    --aws-account-id 111122223333 \  
    --service GLUE_DATA_CATALOG \  
    --authorized-targets arn:aws:sso::111122223333:application/<idc-instance-id>/<lf-  
application-id>
```

The `--authorized-targets` value is the Lake Formation AWS IAM Identity Center application ARN that you obtained in Step 1. This tells Quick which AWS IAM Identity Center application to propagate identity to.

To verify, run:

```
aws quicksight list-identity-propagation-configs \
  --aws-account-id 111122223333
```

Confirm that `GLUE_DATA_CATALOG` appears in the services list.

Step 5: Create a AWS Glue Data Catalog data source

When the author creates a AWS Glue Data Catalog data source, Quick validates accessibility by listing catalogs using the signed-in user's propagated identity. The user sees only the catalogs, databases, and tables that their Lake Formation grants allow.

Verifying the setup

Use the following checklist to verify your configuration:

- The Lake Formation–AWS IAM Identity Center integration returns an `ApplicationArn`.
- Identity propagation scope shows `GLUE_DATA_CATALOG`.
- The IAM role has correct permissions and `sts:SetContext` in the trust policy.
- Data S3 locations are registered with Lake Formation.
- Lake Formation grants are configured for IAM Identity Center users and groups.
- S3 Access Grants are configured on the Athena results bucket (if applicable).
- End-to-end test: sign in as a test user, create a data source, and query a table.

Key considerations and limits

- Permissions are evaluated against the IAM Identity Center end user and their groups.
- The query path runs through Athena. Each user needs S3 Access Grants on the Athena results bucket.
- A data filter has no effect if a full-table `SELECT` is also granted to the same user.
- Best practice: Grant permissions to groups rather than individual users for easier management.
- S3 Access Grants are eventually consistent. Newly added grants may take minutes to propagate.

- Configure fine-grained access control (row, column, and cell filters) in Lake Formation.
- Scope-down policies are evaluated against the end user's identity.

For information about Athena trusted identity propagation, see [Amazon Athena trusted identity propagation](#).

Connecting to Databricks Unity Catalog

Note

Data catalog integration with Databricks Unity Catalog is in Preview for Amazon Quick and is subject to change. Preview features are not typically recommended for critical or production workloads.

The existing Databricks connector in Quick has been enhanced with agentic catalog capabilities and identity enforcement support through OAuth 3LO. Data curators (authors) can use the agentic experience on both new and existing Databricks data sources that were created using a Personal Access Token (PAT).

Setting up the Databricks connection

The Databricks connector supports catalog integration without requiring a new connector. The following connection options are available:

- The agentic experience (**Explore Data**) is available for existing Databricks data sources that use a Personal Access Token (PAT).
- A single Databricks connection provides access to both metadata and data. This differs from AWS Glue Data Catalog, which requires two separate connections.
- Supported authentication options: Personal Access Token (PAT) or OAuth 3LO.

To create a Databricks data source, navigate to **Create Data Source** and select **Databricks**. Choose your preferred authentication method and enter the connection details.

Note

Databricks Metrics Views are not yet supported.

Using the agentic experience with Databricks Unity Catalog

After you connect to Databricks, choose **Explore data** on the data source to launch the agentic experience. The flow is the same as the broader agentic catalog experience:

- **Natural language discovery** – Describe your use case and the AI agent finds relevant tables from Unity Catalog.
- **Bulk create assets** – Create multiple datasets and topics in a single conversational flow.
- **Semantic inheritance** – Table and column descriptions, and primary and foreign key relationships, are automatically inherited from Unity Catalog onto the created datasets.

For details on the full agentic flow, see [The agentic catalog experience](#).

Setting up identity enforcement with Databricks 3LO

Quick supports three-legged OAuth (3LO) with end-user identities for Databricks. When enabled, 3LO enforces per-user data permissions based on Databricks Unity Catalog access control lists (ACLs). Each end user sees only the data that their Databricks permissions allow.

Important

Identity enforcement works only for DirectQuery datasets. If a dataset is switched to SPICE or has transformations applied, identity enforcement does not apply.

Administrator setup: Client application configuration

The Quick account administrator must configure a client application before authors can create 3LO data sources. Complete the following steps:

1. In the Quick console, choose **Manage Account**. Under **Security**, choose **Manage OAuth Client Applications**.
2. Choose **Add OAuth client application** and provide the following information:
 - Config ID
 - OAuth client application name
 - Authentication type
 - Client ID

- Client secret
 - Token endpoint URL
 - Authorization endpoint URL
 - OAuth scopes
 - VPC connection ARN (optional)
 - Data source type (select **DATABRICKS**)
3. Choose **Add** to save the client application.

Author setup: Creating a 3LO data source

After the administrator configures the client application, authors can create a Databricks data source with 3LO authentication.

1. Navigate to the **Create Data Source** page.
2. Select **Databricks**.
3. Choose **3LO** as the authentication method.
4. Enter the required connection information.
5. Enter your username (the email address associated with your Databricks account).
6. Select the option to enforce trusted identity enforcement (TIP). This ensures that permissions defined in Databricks, including table-level access and data filters, are enforced for end users.

Important

You must select this option if end-user identities must be propagated. Without it, all users share the author's Databricks permissions.

7. Choose **Create data source**.

After the data source is created, use the standard Quick workflow or the agentic experience (**Explore Data**) to create datasets, topics, and dashboards.

End-user experience: Enforcing data permissions

To enforce per-user data permissions for dashboard consumers, complete the following steps:

1. The author shares both the data source and the dashboard with the end user.

2. The end user signs in to Quick and navigates to the **Data Source** page.
3. The end user enters their own Databricks credentials (their OAuth login) and saves.
4. When the end user accesses the dashboard, they see only the data that their Databricks permissions allow.

Important considerations

- Identity enforcement works only for DirectQuery datasets.
- The agentic catalog experience (discover, create, inherit) works with both PAT and 3LO authentication. Identity enforcement through 3LO is optional and additive. It is not a prerequisite for the agentic flow.
- If you do not opt for 3LO, you can use PAT and manually manage data permissions by using row-level security (RLS) and column-level security (CLS) rules in Quick.

Data source quotas

Data sources that you use with Amazon Quick Sight must conform to the following quotas.

Topics

- [SPICE quotas for imported data](#)
- [Quotas for direct SQL queries](#)

SPICE quotas for imported data

When you create a new dataset in Amazon Quick Sight, [SPICE](#) limits the number of rows you can add to a dataset. You can ingest data into SPICE from a query or from a file. Each file can have up to 2,000 columns. Each column name can have up to 127 Unicode characters. Each field can have up to 2,047 Unicode characters. If you use the new data preparation experience to create your SPICE dataset, each field can have up to 65,534 Unicode characters.

To retrieve a subset of data from a larger set, you can deselect columns or apply filters to reduce the size of the data. If you are importing from Amazon S3, each manifest can specify up to 1,000 files.

Quotas for SPICE are as follows:

- 2,047 Unicode characters for each field. (65,534 Unicode characters with new data preparation experience)
- 127 Unicode characters for each column name
- 2,000 columns for each file
- 1,000 files for each manifest
- For Standard edition, 25 million (25,000,000) rows or 25 GB for each dataset
- For Enterprise edition, 2 billion (2,000,000,000) rows or 2 TB for each dataset

All quotas apply to SPICE datasets with row-level security, as well.

In rare cases, if you're ingesting large rows into SPICE, you might reach the quota for gigabytes per dataset before you reach the quota on rows. The size is based on the SPICE capacity the data occupies after ingestion into SPICE.

Quotas for direct SQL queries

If you aren't importing data into SPICE, different quotas apply for space and time. For operations such as connecting, sampling data for a dataset, and generating visuals, timeouts can occur. In some cases, these are timeout quotas set by the source database engine. In other cases, such as visualizing, Amazon Quick Sight generates a timeout after 2 minutes.

However, not all database drivers react to the 2-minute timeout, for example Amazon Redshift. In these cases, the query runs for as long as it takes for the response to return, which can result in long-running queries on your database. When this happens, you can cancel the query from the database server to free up database resources. Follow the instructions for your database server about how to do this. For example, for more information on how to cancel queries in Amazon Redshift, see [Canceling a query in Amazon Redshift](#), and [Implementing workload management in Amazon Redshift](#) in the *Amazon Redshift Database Developer Guide*.

Each result set from a direct query can have up to 2,000 columns. Each column name can have up to 127 Unicode characters. If you want to retrieve data from a larger table, you can use one of several methods to reduce the size of the data. You can deselect columns, or apply filters. In a SQL query, you can also use predicates, such as WHERE, HAVING. If your visuals time out during a direct query, you can simplify your query to optimize execution time or you can import the data into SPICE.

Quotas for queries are as follows:

- 127 Unicode characters for each column name.
- 2,000 columns for each dataset.
- 2-minute quota for generating a visual, or an optional dataset sample.
- Data source timeout quotas apply (varies for each database engine).

Supported data types and values

Amazon Quick Sight currently supports the following primitive data types: `Date`, `Decimal`, `Integer`, and `String`. The following data types are supported in SPICE: `Date`, `Decimal-fixed`, `Decimal-float`, `Integer`, and `String`. Quick Sight accepts Boolean values by promoting them to integers. It can also derive geospatial data types. Geospatial data types use metadata to interpret the physical data type. Latitude and longitude are numeric. All other geospatial categories are strings.

Make sure that any table or file that you use as a data source contains only fields that can be implicitly converted to these data types. Amazon Quick Sight skips any fields or columns that can't be converted. If you get an error that says "fields were skipped because they use unsupported data types", alter your query or table to remove or recast unsupported data types.

String and text data

Fields or columns that contain characters are called *strings*. A field with the data type of `STRING` can initially contain almost any type of data. Examples include names, descriptions, phone numbers, account numbers, JSON data, cities, post codes, dates, and numbers that can be used to calculate. These types are sometimes called textual data in a general sense, but not in a technical sense. Quick Sight doesn't support binary and character large objects (BLOBs) in dataset columns. In the Quick Sight documentation, the term "text" always means "string data".

The first time you query or import the data, Quick Sight tries to interpret the data that it identifies as other types, for example dates and numbers. It's a good idea to verify that the data types assigned to your fields or columns are correct.

For each string field in imported data, Quick Sight uses a field length of 8 bytes plus the UTF-8 encoded character length. Amazon Quick Sight supports UTF-8 file encoding, but not UTF-8 (with BOM).

Date and time data

Fields with a data type of Date also include time data, and are also known as Datetime fields. Quick Sight supports dates and times that use [supported date formats](#).

Quick Sight uses UTC time for querying, filtering, and displaying date data. When date data doesn't specify a time zone, Quick Sight assumes UTC values. When date data does specify a time zone, Quick Sight converts it to display in UTC time. For example, a date field with a time zone offset like **2015-11-01T03:00:00-08:00** is converted to UTC and displayed in Amazon Quick Sight as **2015-11-01T15:30:00**.

For each DATE field in imported data, Quick Sight uses a field length of 8 bytes. Quick Sight supports UTF-8 file encoding, but not UTF-8 (with BOM).

Numeric data

Numeric data includes integers and decimals. Integers with a data type of INT are negative or positive numbers that don't have a decimal place. Quick Sight doesn't distinguish between large and small integers. Integers over a value of 9007199254740991 or $2^{53} - 1$ might not display exactly or correctly in a visual.

Decimals with the data type of Decimal are negative or positive numbers that contain at least one decimal place before or after the decimal point. When you choose Direct Query mode, all non-integer decimal types are marked as Decimal and the underlying engine handles the precision of the datapoint based on the data source's supported behaviors. For more information on supported data source types, see [Supported data types and values](#).

When you store your dataset in SPICE, you can choose to store your decimal values as fixed or float decimal types. Decimal-fixed data types use the format of decimal (18, 4) that allow 18 digits total and up to 4 digits after the decimal point. Decimal-fixed data types are a good choice to conduct exact mathematical operations, but Quick Sight rounds the value to the nearest ten thousandth place when the value is ingested into SPICE.

Decimal-float data types provide approximately 16 significant digits of accuracy to a value. The significant digits can be on either side of the decimal point to support numbers with many decimal places and higher numbers at the same time. For example, the Decimal-float data type supports the number 12345.1234567890 or the number 1234567890.12345. If you work with very small numbers that are close to 0, the Decimal-float data type supports up to 15 digits to the right of the decimal point, for example 0.123451234512345. The maximum value that this

data type supports is 1.8×10^{308} to minimize the probability of an overflow error with your data set.

The Decimal-float data type is inexact and some values are stored as approximations instead of the real value. This may result in slight discrepancies when you store and return some specific values. The following considerations apply to the Decimal-float data type.

- If the dataset that you're using comes from an Amazon S3 data source, SPICE assigns the Decimal-float decimal type to all numeric decimal values.
- If the dataset that you're using comes from a database, SPICE uses the decimal type that the value is assigned in the database. For example, if the value is assigned a fixed-point numeric value in the database, the value will be a Decimal-fixed type in SPICE.

For existing SPICE datasets that contain fields that can be converted to the Decimal-float data type, a pop-up appears in the **Edit dataset** page. To convert fields of an existing dataset to the Decimal-float data type, choose **UPDATE FIELDS**. If you don't want to opt in, choose **DO NOT UPDATE FIELDS**. The **Update fields** pop up appears every time you open the **Edit dataset** page until the dataset is saved and published.

Supported data types from external data sources

The following table lists data types that are supported when using the following data sources with Amazon Quick Sight.

Database engine or source	Numeric data types	String data types	Datetime data types	Boolean data types
Amazon Athena, Presto, Starburst, Trino	• bigint • decimal • double • integer • real	• char • varchar	• date • timestamp	• boolean

Database engine or source	Numeric data types	String data types	Datetime data types	Boolean data types
	• smallint • tinyint			
Amazon Aurora, MariaDB, and MySQL	• bigint • decimal • double • int • integer • mediumint • numeric • smallint • tinyint	• char • enum • set • text • varchar	• date • datetime • timestamp • year	
Amazon OpenSearch Service	• byte • integer • long • float • double	• string (keyword string field type in OpenSearch Service) • ip	• timestamp	• boolean • binary

Database engine or source	Numeric data types	String data types	Datetime data types	Boolean data types
Oracle	• bigint • decimal • decimal • int • money • numeric • real • smallint • smallmoney • tinyint	• char • nchar • nvarchar • text • varchar	• date • datetime • datetime2 • datetimeoffset • smalldatetime	bit

Database engine or source	Numeric data types	String data types	Datetime data types	Boolean data types
PostgreSQL	• bigint • decimal • double • integer • numeric • precision • real • smallint	• char • character • text • varchar • varying character	• date • timestamp	• boolean
Apache Spark	• bigint • decimal • double • integer • real • smallint • tinyint	• varchar	• date • timestamp	• boolean

Database engine or source	Numeric data types	String data types	Datetime data types	Boolean data types
Snowflake	• bigint • byteint • decimal • double • doubleprecision • float • float4 • float8 • int • integer • number • numeric • real • smallint • tinyint	• char • character • string • text • varchar	• date • datetime • time • timestamp • timestamp_*	• boolean

Database engine or source	Numeric data types	String data types	Datetime data types	Boolean data types
Microsoft SQL Server	• bigint • bit • decimal • int • money • numeric • real • smallint • smallmoney • tinyint	• char • nchar • nvarchar • text • varchar	• date • datetime • datetime2 • smalldatetime	• bit

Supported date formats

Amazon Quick Sight supports the date and time formats described in this section. Before you add data to Amazon Quick Sight, check if your date format is compatible. If you need to use an unsupported format, see [Using unsupported or custom dates](#).

The supported formats vary depending on the data source type, as follows:

Data source	Clocks	Date formats
File uploads	Both 24-hour and	Supported date and time formats are described in the Joda API documentation.

Data source	Clocks	Date formats
Amazon S3 sources	12-hour clocks	For a complete list of Joda date formats, see Class DateTimeFormat on the Joda website.
Athena		
Salesforce		For datasets stored in memory (SPICE), Amazon Quick Sight supports dates in the following range: Jan 1, 0001 00:00:00 UTC through Dec 31, 9999, 23:59:59 UTC .

Data source	Clocks	Date formats
Relational databases sources	24-hour clock only	The following data and time formats: 1. <code>dd/MM/yyyy HH:mm:ss</code> , for example <code>31/12/2016 15:30:00</code>. 2. <code>dd/MM/yyyy</code> , for example <code>31/12/2016</code>. 3. <code>dd/MMM/yyyy HH:mm:ss</code> , for example <code>31/Dec/2016 15:30:00</code>. 4. <code>dd/MMM/yyyy</code> , for example <code>31/Dec/2016</code>. 5. <code>dd-MMM-yyyy HH:mm:ss</code> , for example <code>31-Dec-2016 15:30:00</code>. 6. <code>dd-MMM-yyyy</code> , for example <code>31-Dec-2016</code>. 7. <code>dd-MM-yyyy HH:mm:ss</code> , for example <code>31-12-2016 15:30:00</code>. 8. <code>dd-MM-yyyy</code> , for example <code>31-12-2016</code>. 9. <code>MM/dd/yyyy HH:mm:ss</code> , for example <code>12/31/2016 15:30:00</code>. 10. <code>MM/dd/yyyy</code> , for example <code>12/31/2016</code>. 11. <code>MM-dd-yyyy HH:mm:ss</code> , for example <code>12-31-2016 15:30:00</code>. 12. <code>MM-dd-yyyy</code> , for example <code>12-31-2016</code>. 13. <code>MMM/dd/yyyy HH:mm:ss</code> , for example <code>Dec/31/2016 15:30:00</code>. 14.

Data source	Clocks	Date formats
		MMM/dd/yyyy , for example Dec/31/2016. 15. MMM-dd-yyyy HH:mm:ss , for example Dec-31-2016 15:30:00. 16. MMM-dd-yyyy , for example Dec-31-2016. 17. yyyy/MM/dd HH:mm:ss , for example 2016/12/31 15:30:00. 18. yyyy/MM/dd , for example 2016/12/31. 19. yyyy/MMM/dd HH:mm:ss , for example 2016/Dec/31 15:30:00. 20. yyyy/MMM/dd , for example 2016/Dec/31. 21. yyyy-MM-dd HH:mm:ss , for example 2016-12-31 15:30:00. 22. yyyy-MM-dd , for example 2016-12-31. 23. yyyy-MMM-dd HH:mm:ss , for example 2016-Dec-31 15:30:00. 24. yyyy-MMM-dd , for example 2016-Dec-31. 25. yyyyMMdd 'T' HHmmss , for example 20161231T153000. 26. yyyy-MM-dd 'T' HH:mm:ss , for example 2016-12-31T15:30:00. 27. yyyyMMdd 'T' HHmmss .SSS , for example 20161231T153000.123.

Data source	Clocks	Date formats
		28. MM/dd/yyyy HH:mm:ss.SSS , for example 12/31/2016 15:30:00.123. 29. dd/MM/yyyy HH:mm:ss.SSS , for example 31/12/2016 15:30:00.123. 30. yyyy/MM/dd HH:mm:ss.SSS , for example 2016/12/31 15:30:00.123. 31. MMM/dd/yyyy HH:mm:ss.SSS , for example Dec/31/2016 15:30:00.123. 32. dd/MMM/yyyy HH:mm:ss.SSS , for example 31/Dec/2016 15:30:00.123. 33. yyyy/MMM/dd HH:mm:ss.SSS , for example 2016/Dec/31 15:30:00.123. 34. yyyy-MM-dd'T'HH:mm:ss.SSS , for example 2016-12-31T15:30:00.123. 35. MM-dd-yyyy HH:mm:ss.SSS , for example 12-31-2016 15:30:00.123. 36. dd-MM-yyyy HH:mm:ss.SSS , for example 31-12-2016 15:30:00.123. 37. yyyy-MM-dd HH:mm:ss.SSS , for example 2016-12-31 15:30:00.123. 38. MMM-dd-yyyy HH:mm:ss.SSS , for example Dec-31-2016 15:30:00.123. 39. dd-MMM-yyyy HH:mm:ss.SSS , for example 31-Dec-2016 15:30:00.123.

Data source	Clocks	Date formats
		40. yyyy-MMM-dd HH:mm:ss.SSS , for example 2016-Dec-31 15:30:00.123.

Unsupported values in data

If a field contains values that don't conform with the data type that Amazon Quick Sight assigns to the field, the rows containing those values are skipped. For example, take the following source data.

Sales ID	Sales Date	Sales Amount
001	10/14/2015	12.43
002	5/3/2012	25.00
003	Unknown	18.17
004	3/8/2009	86.02

Amazon Quick Sight interprets **Sales Date** as a date field and drops the row containing a nondate value, so only the following rows are imported.

Sales ID	Sales Date	Sales Amount
001	10/14/2015	12.43
002	5/3/2012	25.00
004	3/8/2009	86.02

In some cases, a database field might contain values that the JDBC driver can't interpret for the source database engine. In such cases, the uninterpretable values are replaced by null so that the rows can be imported. The only known occurrence of this issue is with MySQL date, datetime, and timestamp fields that have all-zero values, for example **0000-00-00 00:00:00**. For example, take the following source data.

Sales ID	Sales Date	Sales Amount
001	2004-10-12 09:14:27	12.43
002	2012-04-07 12:59:03	25.00

003	0000-00-00 00:00:00	18.17
004	2015-09-30 01:41:19	86.02

In this case, the following data is imported.

Sales ID	Sales Date	Sales Amount
001	2004-10-12 09:14:27	12.43
002	2012-04-07 12:59:03	25.00
003	(null)	18.17
004	2015-09-30 01:41:19	86.02

Working with datasets

Datasets are the foundation of your Quick Sight analytics, serving as the prepared and structured data sources that power your analyses and dashboards. Once you've created datasets from your data sources, you need to manage them effectively throughout their lifecycle to ensure reliable, secure, and collaborative analytics.

This section covers the complete dataset management workflow, from editing and versioning datasets to sharing them with team members and implementing security controls. You'll learn how to maintain dataset integrity while supporting collaborative analytics, track which analyses depend on your datasets, and implement both row-level and column-level security to protect sensitive information. Whether you're preparing datasets for team use, troubleshooting analysis issues, or implementing data governance policies, these topics provide the essential knowledge for effective dataset management in Quick Sight.

Topics

- [Creating datasets](#)
- [Editing datasets](#)
- [Reverting datasets back to previous published versions](#)
- [Duplicating datasets](#)
- [Sharing datasets](#)
- [Tracking dashboards and analyses that use a dataset](#)
- [Using dataset parameters in Amazon Quick](#)
- [Using row-level security in Amazon Quick](#)
- [Using column-level security to restrict access to a dataset](#)

- [Running queries as an IAM role in Amazon Quick](#)
- [Deleting datasets](#)
- [Adding a dataset to an analysis](#)

Creating datasets

You can create datasets from new or existing data sources in Amazon Quick. You can use a variety of database data sources to provide data to Amazon Quick. This includes Amazon RDS instances and Amazon Redshift clusters. It also includes MariaDB, Microsoft SQL Server, MySQL, Oracle, and PostgreSQL instances in your organization, Amazon EC2, or similar environments.

Topics

- [Creating datasets using new data sources](#)
- [Creating a dataset using an existing data source](#)
- [Creating a dataset using an existing dataset in Amazon Quick](#)

Creating datasets using new data sources

When you create a dataset based on an AWS service like Amazon RDS, Amazon Redshift, or Amazon EC2, data transfer charges might apply when consuming data from that source. Those charges might also vary depending on whether that AWS resource is in the home AWS Region that you chose for your Amazon Quick account. For details on pricing, see the pricing page for the service in question.

When creating a new database dataset, you can select one table, join several tables, or create a SQL query to retrieve the data that you want. You can also change whether the dataset uses a direct query or instead stores data in [SPICE](#).

To create a new dataset

1. To create a dataset, choose **New data set** on the **Data** page. You can then create a dataset based on an existing dataset or data source, or connect to a new data source and base the dataset on that.
2. Provide connection information to the data source:
 - For local text or Microsoft Excel files, you can simply identify the file location and upload the file.

- For Amazon S3, provide a manifest identifying the files or buckets that you want to use, and also the import settings for the target files.
- For Amazon Athena, all Athena databases for your AWS account are returned. No additional credentials are required.
- For Salesforce, provide credentials to connect with.
- For Amazon Redshift, Amazon RDS, Amazon EC2, or other database data sources, provide information about the server and database that host the data. Also provide valid credentials for that database instance.

Creating a dataset from a database

The following procedures walk you through connecting to database data sources and creating datasets. To create datasets from AWS data sources that your Amazon Quick account autodiscovered, use [Creating a dataset from an autodiscovered Amazon Redshift cluster or Amazon RDS instance](#). To create datasets from any other database data sources, use [Creating a dataset using a database that's not autodiscovered](#).

Creating a dataset from an autodiscovered Amazon Redshift cluster or Amazon RDS instance

Use the following procedure to create a connection to an autodiscovered AWS data source.

To create a connection to an autodiscovered AWS data source

1. Check [Data source quotas](#) to make sure that your target table or query doesn't exceed data source quotas.
2. Confirm that the database credentials you plan to use have appropriate permissions as described in [Required permissions](#).
3. Make sure that you have configured the cluster or instance for Amazon Quick access by following the instructions in [Network and database configuration requirements](#).
4. On the Amazon Quick start page, choose **Data**.
5. Choose **Create** then choose **New dataset**.
6. Choose either the **RDS** or the **Redshift Auto-discovered** icon, depending on the AWS service that you want to connect to.
7. Enter the connection information for the data source, as follows:
 - For **Data source name**, enter a name for the data source.

- For **Instance ID**, choose the name of the instance or cluster that you want to connect to.
 - **Database name** shows the default database for the **Instance ID** cluster or instance. To use a different database on that cluster or instance, enter its name.
 - For **UserName**, enter the user name of a user account that has permissions to do the following:
 - Access the target database.
 - Read (perform a SELECT statement on) any tables in that database that you want to use.
 - For **Password**, enter the password for the account that you entered.
8. Choose **Validate connection** to verify your connection information is correct.
 9. If the connection validates, choose **Create data source**. If not, correct the connection information and try validating again.

 Note

Amazon Quick automatically secures connections to Amazon RDS instances and Amazon Redshift clusters by using Secure Sockets Layer (SSL). You don't need to do anything to enable this.

10. Choose one of the following:

- **Custom SQL**

On the next screen, you can choose to write a query with the **Use custom SQL** option. Doing this opens a screen named **Enter custom SQL query**, where you can enter a name for your query, and then enter the SQL. For best results, compose the query in a SQL editor, and then paste it into this window. After you name and enter the query, you can choose **Edit/Preview data** or **Confirm query**. Choose **Edit/Preview data** to immediately go to data preparation. Choose **Confirm query** to validate the SQL and make sure that there are no errors.

- **Choose tables**

To connect to specific tables, for **Schema: contain sets of tables**, choose **Select** and then choose a schema. In some cases where there is only a single schema in the database, that schema is automatically chosen, and the schema selection option isn't displayed.

To prepare the data before creating an analysis, choose **Edit/Preview data** to open data preparation. Use this option if you want to join to more tables.

Otherwise, after choosing a table, choose **Select**.

11. Choose one of the following options:

- Prepare the data before creating an analysis. To do this, choose **Edit/Preview data** to open data preparation for the selected table. For more information about data preparation, see [Preparing dataset examples](#).
- Create a dataset and analysis using the table data as-is and to import the dataset data into SPICE for improved performance (recommended). To do this, check the table size and the SPICE indicator to see if you have enough capacity.

If you have enough SPICE capacity, choose **Import to SPICE for quicker analytics**, and then create an analysis by choosing **Visualize**.

 Note

If you want to use SPICE and you don't have enough space, choose **Edit/Preview data**. In data preparation, you can remove fields from the dataset to decrease its size. You can also apply a filter or write a SQL query that reduces the number of rows or columns returned. For more information about data preparation, see [Preparing dataset examples](#).

- To create a dataset and an analysis using the table data as-is, and to have the data queried directly from the database, choose the **Directly query your data** option. Then create an analysis by choosing **Visualize**.

Creating a dataset using a database that's not autodiscovered

Use the following procedure to create a connection to any database other than an autodiscovered Amazon Redshift cluster or Amazon RDS instance. Such databases include Amazon Redshift clusters and Amazon RDS instances that are in a different AWS Region or are associated with a different AWS account. They also include MariaDB, Microsoft SQL Server, MySQL, Oracle, and PostgreSQL instances that are on-premises, in Amazon EC2, or in some other accessible environment.

To create a connection to a database that isn't an autodiscovered Amazon Redshift cluster or RDS instance

1. Check [Data source quotas](#) to make sure that your target table or query doesn't exceed data source quotas.
2. Confirm that the database credentials that you plan to use have appropriate permissions as described in [Required permissions](#).
3. Make sure that you have configured the cluster or instance for Amazon Quick access by following the instructions in [Network and database configuration requirements](#).
4. On the Amazon Quick start page, choose **Manage data**.
5. Choose **Create** then choose **New data set**.
6. Choose the **Redshift Manual connect** icon if you want to connect to an Amazon Redshift cluster in another AWS Region or associated with a different AWS account. Or choose the appropriate database management system icon to connect to an instance of Amazon Aurora, MariaDB, Microsoft SQL Server, MySQL, Oracle, or PostgreSQL.
7. Enter the connection information for the data source, as follows:
 - For **Data source name**, enter a name for the data source.
 - For **Database server**, enter one of the following values:
 - For an Amazon Redshift cluster or Amazon RDS instance, enter the endpoint of the cluster or instance without the port number. For example, if the endpoint value is `clustername.1234abcd.us-west-2.redshift.amazonaws.com:1234`, then enter `clustername.1234abcd.us-west-2.redshift.amazonaws.com`. You can get the endpoint value from the **Endpoint** field on the cluster or instance detail page in the AWS console.
 - For an Amazon EC2 instance of MariaDB, Microsoft SQL Server, MySQL, Oracle, or PostgreSQL, enter the public DNS address. You can get the public DNS value from the **Public DNS** field on the instance detail pane in the Amazon EC2 console.
 - For a non-Amazon EC2 instance of MariaDB, Microsoft SQL Server, MySQL, Oracle, or PostgreSQL, enter the hostname or public IP address of the database server. If you are using Secure Sockets Layer (SSL) for a secured connection (recommended), you likely need to provide the hostname to match the information required by the SSL certificate. For a list of accepted certificates see [Amazon Quick SSL and CA certificates](#).
 - For **Port**, enter the port that the cluster or instance uses for connections.
 - For **Database name**, enter the name of the database that you want to use.

- For **UserName**, enter the user name of a user account that has permissions to do the following:
 - Access the target database.
 - Read (perform a SELECT statement on) any tables in that database that you want to use.
 - For **Password**, enter the password associated with the account you entered.
8. (Optional) If you are connecting to anything other than an Amazon Redshift cluster and you *don't* want a secured connection, make sure that **Enable SSL** is clear. *We strongly recommend leaving this checked*, because an unsecured connection can be open to tampering.

For more information on how the target instance uses SSL to secure connections, see the documentation for the target database management system. Amazon Quick doesn't accept self-signed SSL certificates as valid. For a list of accepted certificates, see [Amazon Quick SSL and CA certificates](#).

Amazon Quick automatically secures connections to Amazon Redshift clusters by using SSL. You don't need to do anything to enable this.

Some databases, such as Presto and Apache Spark, must meet additional requirements before Amazon Quick can connect. For more information, see [Creating a data source using Presto](#), or [Creating a data source using Apache Spark](#).

9. (Optional) Choose **Validate connection** to verify your connection information is correct.
10. If the connection validates, choose **Create data source**. If not, correct the connection information and try validating again.
11. Choose one of the following:
- **Custom SQL**

On the next screen, you can choose to write a query with the **Use custom SQL** option. Doing this opens a screen named **Enter custom SQL query**, where you can enter a name for your query, and then enter the SQL. For best results, compose the query in a SQL editor, and then paste it into this window. After you name and enter the query, you can choose **Edit/Preview data** or **Confirm query**. Choose **Edit/Preview data** to immediately go to data preparation. Choose **Confirm query** to validate the SQL and make sure that there are no errors.

- **Choose tables**

To connect to specific tables, for **Schema: contain sets of tables**, choose **Select** and then choose a schema. In some cases where there is only a single schema in the database, that schema is automatically chosen, and the schema selection option isn't displayed.

To prepare the data before creating an analysis, choose **Edit/Preview data** to open data preparation. Use this option if you want to join to more tables.

Otherwise, after choosing a table, choose **Select**.

12. Choose one of the following options:

- Prepare the data before creating an analysis. To do this, choose **Edit/Preview data** to open data preparation for the selected table. For more information about data preparation, see [Preparing dataset examples](#).
- Create a dataset and an analysis using the table data as-is and import the dataset data into SPICE for improved performance (recommended). To do this, check the table size and the SPICE indicator to see if you have enough space.

If you have enough SPICE capacity, choose **Import to SPICE for quicker analytics**, and then create an analysis by choosing **Visualize**.

 Note

If you want to use SPICE and you don't have enough space, choose **Edit/Preview data**. In data preparation, you can remove fields from the dataset to decrease its size. You can also apply a filter or write a SQL query that reduces the number of rows or columns returned. For more information about data preparation, see [Preparing dataset examples](#).

- Create a dataset and an analysis using the table data as-is and have the data queried directly from the database. To do this, choose the **Directly query your data** option. Then create an analysis by choosing **Visualize**.

Creating a dataset using an existing data source

After you make an initial connection to a Salesforce, AWS data store, or other database data source, Amazon Quick saves the connection information. It adds the data source to the **FROM**

EXISTING DATA SOURCES section of the **Create a Data Set** page. You can use these existing data sources to create new datasets without respecifying connection information.

Creating a dataset using an existing Amazon S3 data source

Use the following procedure to create a dataset using an existing Amazon S3 data source.

To create a dataset using an existing S3 data source

1. On the Amazon Quick start page, choose **Data**.
2. Choose **Create** then choose **New dataset**.
3. Choose the Amazon S3 data source to use.
4. To prepare the data before creating the dataset, choose **Edit/Preview data**. To create an analysis using the data as-is, choose **Visualize**.

Creating a dataset using an existing Amazon Athena data source

To create a dataset using an existing Amazon Athena data source, use the following procedure.

To create a dataset from an existing Athena connection profile

1. On the Amazon Quick start page, choose **Data**.
2. Choose **Create** then choose **New data set**.

Choose the connection profile icon for the existing data source that you want to use.

Connection profiles are labeled with the data source icon and the name provided by the person who created the connection.

3. Choose **Create data set**.

Amazon Quick creates a connection profile for this data source based only on the Athena workgroup. The database and table aren't saved.

4. On the **Choose your table** screen, do one of the following:
 - To write a SQL query, choose **Use custom SQL**.
 - To choose a database and table, first select your database from the **Database** list. Next, choose a table from the list that appears for your database.

Create a dataset using an existing Salesforce data source

Use the following procedure to create a dataset using an existing Salesforce data source.

To create a dataset using an existing Salesforce data source

1. On the Amazon Quick start page, choose **Data**.
2. Choose **Create** then choose **New data set**.
3. Choose the Salesforce data source to use.
4. Choose **Create Data Set**.
5. Choose one of the following:

- **Custom SQL**

On the next screen, you can choose to write a query with the **Use custom SQL** option. Doing this opens a screen named **Enter custom SQL query**, where you can enter a name for your query, and then enter the SQL. For best results, compose the query in a SQL editor, and then paste it into this window. After you name and enter the query, you can choose **Edit/Preview data** or **Confirm query**. Choose **Edit/Preview data** to immediately go to data preparation. Choose **Confirm query** to validate the SQL and make sure that there are no errors.

- **Choose tables**

To connect to specific tables, for **Data elements: contain your data**, choose **Select** and then choose either **REPORT** or **OBJECT**.

To prepare the data before creating an analysis, choose **Edit/Preview data** to open data preparation. Use this option if you want to join to more tables.

Otherwise, after choosing a table, choose **Select**.

6. On the next screen, choose one of the following options:
 - To create a dataset and an analysis using the data as-is, choose **Visualize**.

Note

If you don't have enough [SPICE](#) capacity, choose **Edit/Preview data**. In data preparation, you can remove fields from the dataset to decrease its size or apply a

filter that reduces the number of rows returned. For more information about data preparation, see [Preparing dataset examples](#).

- To prepare the data before creating an analysis, choose **Edit/Preview data** to open data preparation for the selected report or object. For more information about data preparation, see [Preparing dataset examples](#).

Creating a dataset using an existing database data source

Use the following procedure to create a dataset using an existing database data source.

To create a dataset using an existing database data source

1. On the Amazon Quick start page, choose **Data**.
2. Choose **Create** then choose **New data set**.
3. Choose the database data source to use, and then choose **Create Data Set**.
4. Choose one of the following:

- **Custom SQL**

On the next screen, you can choose to write a query with the **Use custom SQL** option. Doing this opens a screen named **Enter custom SQL query**, where you can enter a name for your query, and then enter the SQL. For best results, compose the query in a SQL editor, and then paste it into this window. After you name and enter the query, you can choose **Edit/Preview data** or **Confirm query**. Choose **Edit/Preview data** to immediately go to data preparation. Choose **Confirm query** to validate the SQL and make sure that there are no errors.

- **Choose tables**

To connect to specific tables, for **Schema: contain sets of tables**, choose **Select** and then choose a schema. In some cases where there is only a single schema in the database, that schema is automatically chosen, and the schema selection option isn't displayed.

To prepare the data before creating an analysis, choose **Edit/Preview data** to open data preparation. Use this option if you want to join to more tables.

Otherwise, after choosing a table, choose **Select**.

5. Choose one of the following options:

- Prepare the data before creating an analysis. To do this, choose **Edit/Preview data** to open data preparation for the selected table. For more information about data preparation, see [Preparing dataset examples](#).
- Create a dataset and an analysis using the table data as-is and import the dataset data into [SPICE](#) for improved performance (recommended). To do this, check the SPICE indicator to see if you have enough space.

If you have enough SPICE capacity, choose **Import to SPICE for quicker analytics**, and then create an analysis by choosing **Visualize**.

 Note

If you want to use SPICE and you don't have enough space, choose **Edit/Preview data**. In data preparation, you can remove fields from the dataset to decrease its size. You can also apply a filter or write a SQL query that reduces the number of rows or columns returned. For more information about data preparation, see [Preparing dataset examples](#).

- Create a dataset and an analysis using the table data as-is and have the data queried directly from the database. To do this, choose the **Directly query your data** option. Then create an analysis by choosing **Visualize**.

Creating a dataset using an existing dataset in Amazon Quick

After you create a dataset in Amazon Quick, you can create additional datasets using it as a source. When you do this, any data preparation that the parent dataset contains, such as any joins or calculated fields, is kept. You can add additional preparation to the data in the new child datasets, such as joining new data and filtering data. You can also set up your own data refresh schedule for the child dataset and track the dashboards and analyses that use it.

Child datasets that are created using a dataset with RLS rules active as a source inherit the parent dataset's RLS rules. Users who are creating a child dataset from a larger parent dataset can only see the data that they have access to in the parent dataset. Then, you can add more RLS rules to the new child dataset in addition to the inherited RLS rules to further manage who can access the data that is in the new dataset. You can only create child datasets from datasets with RLS rules active in Direct Query.

Creating datasets from existing Quick datasets has the following advantages:

- **Central management of datasets** – Data engineers can easily scale to the needs of multiple teams within their organization. To do this, they can develop and maintain a few general-purpose datasets that describe the organization's main data models.
- **Reduction of data source management** – Business analysts (BAs) often spend lots of time and effort requesting access to databases, managing database credentials, finding the right tables, and managing Quick data refresh schedules. Building new datasets from existing datasets means that BAs don't have to start from scratch with raw data from databases. They can start with curated data.
- **Predefined key metrics** – By creating datasets from existing datasets, data engineers can centrally define and maintain critical data definitions across their company's many organizations. Examples might be sales growth and net marginal return. With this feature, data engineers can also distribute changes to those definitions. This approach means that their business analysts can get started with visualizing the right data more quickly and reliably.
- **Flexibility to customize data** – By creating datasets from existing datasets, business analysts get more flexibility to customize datasets for their own business needs. They can avoid worry about disrupting data for other teams.

For example, let's say that you're part of an ecommerce central team of five data engineers. You and your team has access to sales, orders, cancellations, and returns data in a database. You have created a Quick dataset by joining 18 other dimension tables through a schema. A key metric that your team has created is the calculated field, order product sales (OPS). Its definition is: $OPS = \text{product quantity} \times \text{price}$.

Your team serves over 100 business analysts across 10 different teams in eight countries. These are the Coupons team, the Outbound Marketing team, the Mobile Platform team, and the Recommendations team. All of these teams use the OPS metric as a base to analyze their own business line.

Rather than manually creating and maintaining hundreds of unconnected datasets, your team reuses datasets to create multiple levels of datasets for teams across the organization. Doing this centralizes data management and allows each team to customize the data for their own needs. At the same time, this syncs updates to the data, such as updates to metric definitions, and maintains row-level and column-level security. For example, individual teams in your organization can use the centralized datasets. They can then combine them with the data specific to their team to create new datasets and build analyses on top of them.

Along with using the key OPS metric, other teams in your organization can reuse column metadata from the centralized datasets that you created. For example, the Data Engineering team can define metadata, such as *name*, *description*, *data type*, and *folders*, in a centralized dataset. All subsequent teams can use it.

Note

Amazon Quick supports creating up to two additional levels of datasets from a single dataset.

For example, from a parent dataset, you can create a child dataset and then a grandchild dataset for a total of three dataset levels.

Creating a dataset from an existing dataset

Use the following procedure to create a dataset from an existing dataset.

To create a dataset from an existing dataset

1. From the Quick start page, choose **Data** in the pane at left.
2. Choose **Create** then choose the dataset that you want to use to create a new dataset.
3. On the page that opens for that dataset, choose the drop-down menu for **Use in analysis**, and then choose **Use in dataset**.

The data preparation page opens and preloads everything from the parent dataset, including calculated fields, joins, and security settings.

4. On the data preparation page that opens, for **Query mode** at bottom left, choose how you want the dataset to pull in changes and updates from the original, parent dataset. You can choose the following options:
 - **Direct query** – This is the default query mode. If you choose this option, the data for this dataset automatically refreshes when you open an associated dataset, analysis, or dashboard. However, the following limitations apply:
 - If the parent dataset allows direct querying, you can use direct query mode in the child dataset.
 - If you have multiple parent datasets in a join, you can choose direct query mode for your child dataset only if all the parents are from the same underlying data source. For example, the same Amazon Redshift connection.

- Direct query is supported for a single SPICE parent dataset. It is not supported for multiple SPICE parent datasets in a join.
 - **SPICE** – If you choose this option, you can set up a schedule for your new dataset to sync with the parent dataset. For more information about creating SPICE refresh schedules for datasets, see [Refreshing SPICE data](#).
5. (Optional) Prepare your data for analysis. For more information about preparing data, see [Preparing data in Amazon Quick Sight](#).
 6. (Optional) Set up row-level or column-level security (RLS/CLS) to restrict access to the dataset. For more information about setting up RLS, see [Using row-level security with user-based rules to restrict access to a dataset](#). For more information about setting up CLS, see [Using column-level security to restrict access to a dataset](#).
-  **Note**

You can set up RLS/CLS on child datasets only. RLS/CLS on parent datasets is not supported.
7. When you're finished, choose **Save & publish** to save your changes and publish the new child dataset. Or choose **Publish & visualize** to publish the new child dataset and begin visualizing your data.

Restricting others from creating new datasets from your dataset

When you create a dataset in Amazon Quick, you can prevent others from using it as a source for other datasets. You can specify if others can use it to create any datasets at all. Or you can specify the type of datasets others can or can't create from your dataset, such as direct query datasets or SPICE datasets.

Use the following procedure to learn how to restrict others from creating new datasets from your dataset.

To restrict others from creating new datasets from your dataset

1. From the Quick start page, choose **Data** in the pane at left.
2. Choose **Create** then choose the dataset that you want to restrict creating new datasets from.
3. On the page that opens for that dataset, choose **Edit dataset**.

4. On the data preparation page that opens, choose **Manage** at upper right, and then choose **Properties**.
5. In the **Dataset properties** pane that opens, choose from the following options:
 - To restrict anyone from creating any type of new datasets from this dataset, turn off **Allow new datasets to be created from this one**.

The toggle is blue when creating new datasets is allowed. It's gray when creating new datasets isn't allowed.

- To restrict others from creating direct query datasets, clear **Allow direct query**.
- To restrict others from creating SPICE copies of your dataset, clear **Allow SPICE copies**.

For more information about SPICE datasets, see [Importing data into SPICE](#).

6. Close the pane.

Editing datasets

You can edit an existing dataset to perform data preparation. For more information about Quick Sight data preparation functionality, see [Preparing data in Amazon Quick Sight](#).

You can open a dataset for editing from the **Datasets** page, or from the analysis page. Editing a dataset from either location modifies the dataset for all analyses that use it.

Things to consider when editing datasets

In two situations, changes to a dataset might cause concern. One is if you deliberately edit the dataset. The other is if your data source has changed so much that it affects the analyses based on it.

Important

Analyses that are in production usage should be protected so they continue to function correctly.

We recommend the following when you're dealing with data changes:

- Document your data sources and datasets, and the visuals that rely upon them. Documentation should include screenshots, fields used, placement in field wells, filters, sorts, calculations, colors,

formatting, and so on. Record everything that you need to recreate the visual. You can also track which Quick Sight resources use a dataset in the dataset management options. For more information, see [Tracking dashboards and analyses that use a dataset](#).

- When you edit a dataset, try not to make changes that might break existing visuals. For example, don't remove columns that are being used in a visual. If you must remove a column, create a calculated column in its place. The replacement column should have the same name and data type as the original.
- If your data source or dataset changes in your source database, adapt your visual to accommodate the change, as described previously. Or you can try to adapt the source database. For example, you might create a view of the source table (document). Then if the table changes, you can adjust the view to include or exclude columns (attributes), change data types, fill null values, and so on. Or, in another circumstance, if your dataset is based on a slow SQL query, you might create a table to hold the results of the query.

If you can't sufficiently adapt the source of the data, recreate the visuals based on your documentation of the analysis.

- If you no longer have access to a data source, your analyses based on that source are empty. The visuals that you created still exist, but they can't display until they have some data to show. This result can happen if permissions are changed by your administrator.
- If you remove the dataset a visual is based on, you might need to recreate it from your documentation. You can edit the visual and select a new dataset to use with it. If you need to consistently use a new file to replace an older one, store your data in a location that is consistently available. For example, you might store your .csv file in Amazon S3 and create an S3 dataset to use for your visuals. For more information on access files stored in S3, see [Creating a dataset using Amazon S3 files](#).

Or you can import the data into a table, and base your visual on a query. This way, the data structures don't change, even if the data contained in them changes.

- To centralize data management, consider creating general, multiple-purpose datasets that others can use to create their own datasets from. For more information, see [Creating a dataset using an existing dataset in Amazon Quick](#).

Editing a dataset from the Datasets page

1. From the Quick start page, choose **Data** at left.

2. On the **Data** page that opens, choose the dataset that you want to edit, and then choose **Edit dataset** at upper right.

The data preparation page opens. For more information about the types of edits you can make to datasets, see [Preparing data in Amazon Quick Sight](#).

Editing a dataset in an analysis

Use the following procedure to edit a dataset from the analysis page.

To edit a dataset from the analysis page

1. In your analysis, choose the pencil icon at the top of the **Fields list** pane.
2. In **Data sets in this analysis** page that opens, choose the three dots at right of the dataset that you want to edit, and then choose **Edit**.

The dataset opens in the data preparation page. For more information about the types of edits you can make to datasets, see [Preparing data in Amazon Quick Sight](#).

Reverting datasets back to previous published versions

When you save and publish changes to a dataset in Amazon Quick Sight, a new version of the dataset is created. At any time, you can see a list of all the previous published versions of that dataset. You can also preview a specific version in that history, or even revert the dataset back to a previous version, if needed.

The following limitations apply to dataset versioning:

- Only the most recent 1,000 versions of a dataset are shown in the publishing history, and are available for versioning.
- After you exceed 1,000 published versions, the oldest versions are automatically removed from the publishing history, and the dataset can no longer be reverted back to them.

Use the following procedure to revert a dataset to a previous published version.

To revert a dataset to a previous published version

1. From the Quick start page, choose **Data**.

2. On the **Data** page, choose a dataset, and then choose **Edit dataset** at upper right.

For more information about editing datasets, see [Editing datasets](#).

3. On the dataset preparation page that opens, choose the **Manage** icon in the blue toolbar at upper right, and then choose **Publishing history**.

A list of previous published versions appears at right.

4. In the **Publishing history** pane, find the version that you want and choose **Revert**.

To preview the version before reverting, choose **Preview**.

The dataset is reverted and a confirmation message appears. The **Publishing history** pane also updates to show the active version of the dataset.

Troubleshooting reverting versions

Sometimes, the dataset can't be reverted to a specific version for one the following reasons:

- The dataset uses one or more data sources that were deleted.

If this error occurs, you can't revert the dataset to a previous version.

- Reverting would make a calculated field not valid.

If this error occurs, you can edit or remove the calculated field, and then save the dataset. Doing this creates a new version of the dataset.

- One or more columns are missing in the data source.

If this error occurs, Quick Sight shows the latest schema from the data source in the preview to reconcile differences between versions. Any calculated field, field name, field type, and filter changes shown in the schema preview are from the version that you want to revert to. You can save this reconciled schema as a new version of the dataset. Or you can return to the active (latest) version by choosing **Preview** on the top (latest) version in the publishing history.

Duplicating datasets

You can duplicate an existing dataset to save a copy of it with a new name. The new dataset is a completely separate copy.

The **Duplicate dataset** option is available if both of the following are true: you own the dataset and you have permission to the data source.

To duplicate a dataset

1. From the Quick start page, choose **Data** at left.
2. Choose the dataset that you want to duplicate.
3. On the dataset details page that opens, choose the drop-down for **Edit dataset**, and then choose **Duplicate**.
4. On the Duplicate dataset page that opens, give the duplicated dataset a name, and then choose **Duplicate**.

The duplicated dataset details page opens. From this page, you can edit the dataset, set up a refresh schedule, and more.

Sharing datasets

You can give other Quick Sight users and groups access to a dataset by sharing it with them. Then they can create analyses from it. If you make them co-owners, they can also refresh, edit, delete, or reshare the dataset.

Sharing a dataset

If you have owner permissions on a dataset, use the following procedure to share it.

To share a dataset

1. From the Quick start page, choose **Data** at left.
2. On the **Data** page, choose the dataset that you want to share.
3. On the dataset details page that opens, choose the **Permissions** tab, and then choose **Add users & groups**.
4. Enter the user or group that you want to share this dataset with, and then choose **Add**. You can only invite users who belong to the same Quick account.

Repeat this step until you have entered information for everyone you want to share the dataset with.

5. For the **Permissions** column, choose a role for each user or group to give them permissions on the dataset.

Choose **User** to allow them to create analyses and datasets from the dataset. Choose **Owner** to allow them to do that and also refresh, edit, delete, and reshare the dataset.

Users receive emails with a link to the dataset. Groups don't receive invitation emails.

Viewing and editing the permissions of users that a dataset is shared with

If you have owner permissions on a dataset, you can use the following procedure to view, edit, or change user access to it.

To view, edit, or change user access to a dataset if you have owner permissions for it

1. From the Quick start page, choose **Data** at left.
2. On the **Data** page, choose the dataset that you want to share.
3. On the dataset details page that opens, choose the **Permissions** tab.

A list of all users and groups with access to the dataset is displayed.

4. (Optional) To change permission roles for a user or group, choose the drop-down menu in the **Permissions** column for the user or group. Then choose either **User** or **Owner**.

Revoking access to a dataset

If you have owner permissions on a dataset, you can use the following procedure to revoke user access to a dataset.

To revoke user access to a dataset if you have owner permissions for it

1. From the Quick start page, choose **Data** at left.
2. On the **Data** page, choose the dataset that you want to share.
3. On the dataset details page that opens, choose the **Permissions** tab.

A list of all users and groups with access to the dataset is displayed.

4. In the **Actions** column for the user or group, choose **Revoke access**.

Tracking dashboards and analyses that use a dataset

When you create a dataset in Quick Sight, you can track which dashboards and analyses use that dataset. This approach is useful when you want to see which resources will be affected when you make changes to a dataset, or want to delete a dataset.

Use the following procedure to see which dashboards and analyses use a dataset.

To track resources that use a dataset

1. From the Quick start page, choose **Data** in the pane at left.
2. On the **Data** page, choose the dataset that you want to track resources for.
3. In the page that opens for that dataset, choose **Edit dataset**.
4. In the data preparation page that opens, choose **Manage** at upper right, and then choose **Usage**.
5. The dashboards and analyses that use the dataset are listed in the pane that opens.

Using dataset parameters in Amazon Quick

In Amazon Quick, authors can use dataset parameters in direct query to dynamically customize their datasets and apply reusable logic to their datasets. A *dataset parameter* is a parameter created at the dataset level. It's consumed by an analysis parameter through controls, calculated fields, filters, actions, URLs, titles, and descriptions. For more information on analysis parameters, see [Parameters in Amazon Quick](#). The following list describes three actions that can be performed with dataset parameters:

- **Custom SQL in direct query** – Dataset owners can insert dataset parameters into the custom SQL of a direct query dataset. When these parameters are applied to a filter control in a Quick analysis, users can filter their custom data faster and more efficiently.
- **Repeatable variables** – Static values that appear in multiple locations in the dataset page can be modified in one action using custom dataset parameters.
- **Move calculated fields to datasets** – Quick authors can copy calculated fields with parameters in an analysis and migrate them to the dataset level. This protects calculated fields at the analysis level from being accidentally modified and calculated fields be shared across multiple analyses.

In some situations, dataset parameters improve filter control performance for direct query datasets that require complex custom SQL and simplify business logic at the dataset level.

Topics

- [Dataset parameter limitations](#)
- [Creating dataset parameters in Amazon Quick](#)
- [Inserting dataset parameters into custom SQL](#)
- [Adding dataset parameters to calculated fields](#)
- [Adding dataset parameters to filters](#)
- [Using dataset parameters in Quick analyses](#)
- [Advanced use cases of dataset parameters](#)

Dataset parameter limitations

This section covers known limitations that you might encounter when working with dataset parameters in Amazon Quick.

- When dashboard readers schedule emailed reports, selected controls don't propagate to the dataset parameters that are included in the report that's attached to the email. Instead, the default values of the parameters are used.
- Dataset parameters can't be inserted into custom SQL of datasets stored in SPICE.
- Dynamic defaults can only be configured on the analysis page of the analysis that is using the dataset. You can't configure a dynamic default at the dataset level.
- The **Select all** option is not supported on multivalue controls of analysis parameters that are mapped to dataset parameters.
- Cascading controls are not supported for dataset parameters.
- Dataset parameters can only be used by dataset filters when the dataset is using direct query.
- In a custom SQL query, only 128 dataset parameters can be used.

Creating dataset parameters in Amazon Quick

Use the following procedures to get started using dataset parameters.

To create a new dataset parameter

1. From the Quick start page, choose **Data** on the left, choose the ellipsis (three dots) next to the dataset that you want to change, and then choose **Edit**.

2. On the **Dataset** page that opens, choose **Parameters** on the left, and then choose the (+) icon to create a new dataset parameter.
3. In the **Create new parameter** pop-up that appears, enter a parameter name in the **Name** box.
4. In the **Data type** dropdown, choose the parameter data type that you want. Supported data types are String, Integer, Number, and Datetime. This option can't be changed after the parameter is created.
5. For **Default value**, enter the default value that you want the parameter to have.

 Note

When you map a dataset parameter to an analysis parameter, a different default value can be chosen. When this happens, the default value configured here is overridden by the new default value.

6. For **Values**, choose the value type that you want the parameter to have. **Single value** parameters support single-select dropdowns, text field, and list controls. **Multiple values** parameters support multi-select dropdown controls. This option can't be changed after the parameter is created.
7. When you are finished configuring the new parameter, choose **Create** to create the parameter.

Inserting dataset parameters into custom SQL

You can insert dataset parameters into the custom SQL of a dataset in direct query mode by referencing it with `<<$parameter_name>>` in the SQL statement. At runtime, dashboard users can enter filter control values that are associated with a dataset parameter. Then, they can see the results in the dashboard visuals after the values propagate to the SQL query. You can use parameters to create basic filters based on customer input in where clauses. Alternatively, you could add `case when` or `if else` clauses to dynamically change the logic of the SQL query based on a parameter's input.

For example, say you want to add a `WHERE` clause to your custom SQL that filters data based on an end user's Region name. In this case, you create a single value parameter called `RegionName`:

```
SELECT *  
FROM transactions  
WHERE region = <<$RegionName>>
```


You can also let users provide multiple values to the parameter:

```
SELECT *
FROM transactions
WHERE region in (<<$RegionNames>>)
```

In the following more complex example, a dataset author refers to two dataset parameters twice based on a user's first and last names that can be selected in a dashboard filter control:

```
SELECT Region, Country, OrderDate, Sales
FROM transactions
WHERE region=
(Case
WHEN <<$UserFIRSTNAME>> In
    (select firstname from user where region='region1')
    and <<$UserLASTNAME>> In
    (select lastname from user where region='region1')
    THEN 'region1'
WHEN <<$UserFIRSTNAME>> In
    (select firstname from user where region='region2')
    and <<$UserLASTNAME>> In
    (select lastname from user where region='region2')
    THEN 'region2'
ELSE 'region3'
END)
```

You can also use parameters in SELECT clauses to create new columns in a dataset from user input:

```
SELECT Region, Country, date,
    (case
    WHEN <<$RegionName>>='EU'
    THEN sum(sales) * 0.93    --convert US dollar to euro
    WHEN <<$RegionName>>='CAN'
    THEN sum(sales) * 0.78    --convert US dollar to Canadian Dollar
    ELSE sum(sales) -- US dollar
    END
    ) as "Sales"
FROM transactions
WHERE region = <<$RegionName>>
```

To create a custom SQL query or to edit an existing query before adding a dataset parameter, see [Using SQL to customize data](#).

When you apply custom SQL with a dataset parameter, <<\$parameter_name>> is used as a placeholder value. When a user chooses one of the parameter values from a control, Quick replaces the placeholder with the values that the user selects on the dashboard.

In the following example, the user enters a new custom SQL query that filters data by state:

```
select * from all_flights
where origin_state_abr = <<$State>>
```

The default value of the parameter is applied to the SQL query and the results appear in the **Preview pane**.

Adding dataset parameters to calculated fields

You can also add dataset parameters to calculated field expressions using the format `${parameter_name}`.

When you create a calculation, you can choose from the existing parameters from the list of parameters under the **Parameters** list. You can't create a calculated field that contains a multivalued parameter.

For more information on adding calculated fields, see [Using calculated fields with parameters in Amazon Quick](#).

Adding dataset parameters to filters

For datasets in direct query mode, dataset authors can use dataset parameters in filters without custom SQL. Dataset parameters can't be added to filters if the dataset is in SPICE.

To add a dataset parameter to a filter

1. Open the dataset page of the dataset that you want to create a filter for. Choose **Filters** on the left, and then choose **Add filter**.
2. Enter the name that you want the filter to have and choose the field that you want filtered in the dropdown.
3. After you create the new filter, navigate to the filter in the **Filters** pane, choose the ellipsis (three dots) next to the filter, and then choose **Edit**.
4. For **Filter type**, choose **Custom filter**.
5. For **Filter condition**, choose the condition that you want.

6. Select the **Use parameter** box and choose the dataset parameter that you want the filter to use.
7. When you are finished making changes, choose **Apply**.

Using dataset parameters in Quick analyses

Once you create a dataset parameter, after you add the dataset to an analysis, map the dataset parameter to a new or existing analysis parameter. After you map a dataset parameter to an analysis parameter, you can use them with filters, controls, and any other analysis parameter feature.

You can manage your dataset parameters in the **Parameters** pane of the analysis that is using the dataset that the parameters belong to. In the **Dataset Parameters** section of the **Parameters** pane, you can choose to see only the unmapped dataset parameters (default). Alternatively, you can choose to see all mapped and unmapped dataset parameters by choosing **ALL** from the **Viewing** dropdown.

Mapping dataset parameters in new Quick analyses

When you create a new analysis from a dataset that contains parameters, you need to map the dataset parameters to the analysis before you can use them. This is also true when you add a dataset with parameters to an analysis. You can view all unmapped parameters in an analysis in the **Parameters** pane of the analysis. Alternatively, choose **VIEW** in the notification message that appears in the top right of the page when you create the analysis or add the dataset.

To map a dataset parameter to an analysis parameter

1. Open the [Quick console](#).
2. Choose the analysis that you want to change.
3. Choose the **Parameters** icon to open the **Parameters** pane.
4. Choose the ellipsis (three dots) next to the dataset parameter that you want to map, choose **Map Parameter**, and then choose the analysis parameter that you want to map your dataset parameter to.

If your analysis doesn't have any analysis parameters, you can choose **Map parameter** and **Create new** to create an analysis parameter that is automatically mapped to the dataset parameter upon creation.

- a. (Optional) In the **Create new parameter** pop-up that appears, for **Name**, enter a name for the new analysis parameter.
- b. (Optional) For **Static default value**, choose the static default value that you want the parameter to have.
- c. (Optional) Choose **Set a dynamic default** to set a dynamic default for the new parameter.
- d. In the **Mapped dataset parameters** table, you will see the dataset parameter that you are mapping to the new analysis parameter. You can add other dataset parameters to this analysis parameter by choosing the **ADD DATASET PARAMETER** dropdown and then choosing the parameter that you want to map. You can unmap a dataset parameter by choosing the **Remove** button next to the dataset parameter that you want to remove.

For more information on creating analysis parameters, see [Setting up parameters in Amazon Quick](#).

When you map a dataset parameter to an analysis parameter, the analysis parameter represents the dataset parameter wherever it is used in the analysis.

You can also map and unmap dataset parameters to analysis parameters in the **Edit parameter** window. To open the **Edit parameter** window, navigate to the **Parameters** pane, choose the ellipsis (three dots) next to the analysis parameter that you want to change, and then choose **Edit parameter**. You can add other dataset parameters to this analysis parameter by choosing the **ADD DATASET PARAMETER** dropdown and then choosing the parameter that you want to map. You can unmap a dataset parameter by choosing the **Remove** button next to the dataset parameter that you want to remove. You can also remove all mapped dataset parameters by choosing **REMOVE ALL**. When you are done making changes, choose **Update**.

When you delete an analysis parameter, all dataset parameters are unmapped from the analysis and appear in the **UNMAPPED** section of the **Parameters** pane. You can only map a dataset parameter to one analysis parameter at a time. To map a dataset parameter to a different analysis parameter, unmap the dataset parameter and then map it to the new analysis parameter.

Adding filter controls to mapped analysis parameters

After you map a dataset parameter to an analysis parameter in Quick, you can create filter controls for filters, actions, calculated fields, titles, descriptions, and URLs.

To add a control to a mapped parameter

1. In the **Parameters** pane of the analysis page, choose the ellipsis (three dots) next to the mapped analysis parameter that you want, and then choose **Add control**.
2. In the **Add control** window that appears, enter the **Name** that you want and choose the **Style** that you want the control to have. For single value controls, choose between Dropdown, List, and Text field. For multivalue controls, choose Dropdown.
3. Choose **Add** to create the control.

Advanced use cases of dataset parameters

This section covers more advanced options and use cases working with dataset parameters and dropdown controls. Use the following walkthroughs to create dynamic dropdown values with dataset parameters.

Using multivalue controls with dataset parameters

When you use dataset parameters that are inserted into the custom SQL of a dataset, the dataset parameters commonly filter data by values from a specific column. If you create a dropdown control and assign the parameter as the value, the dropdown only shows the value that the parameter filtered. The following procedure shows how you can create a control that is mapped to a dataset parameter and shows all unfiltered values.

To populate all assigned values in a dropdown control

1. Create a new single-column dataset in SPICE or direct query that includes all unique values from the original dataset. For example, let's say that your original dataset is using the following custom SQL:

```
select * from all_flights
      where origin_state_abr = <<$State>>
```

To create a single-column table with all unique origin states, apply the following custom SQL to the new dataset:

```
SELECT distinct origin_state_abr FROM all_flights
      order by origin_state_abr asc
```

The SQL expression returns all unique states in alphabetic order. The new dataset does not have any dataset parameters.

2. Enter a **Name** for the new dataset, and then save and publish the dataset. In our example, the new dataset is called `State Codes`.
3. Open the analysis that contains the original dataset, and add the new dataset to the analysis. For information on adding datasets to an existing analysis, see [Adding a dataset to an analysis](#).
4. Navigate to the **Controls** pane and find the dropdown control that you want to edit. Choose the ellipsis (three dots) next to the control, and then choose **Edit**.
5. In the **Format control** that appears on the left, and choose **Link to a dataset field** in the **Values** section.
6. For the **Dataset** dropdown that appears, choose the new dataset that you created. In our example, the `State Codes` dataset is chosen.
7. For the **Field** dropdown that appears, choose the appropriate field. In our example, the `origin_state_abr` field is chosen.

After you finish linking the control to the new dataset, all unique values appear in the control's dropdown. These include the values that are filtered out by the dataset parameter.

Using controls with Select all options

By default, when one or more dataset parameters are mapped to an analysis parameter and added to a control, the `Select all` option is not available. The following procedure shows a workaround that uses the same example scenario from the previous section.

Note

This walkthrough is for datasets that are small enough to load in direct query. If you have a large dataset and want to use the `Select All` option, it is recommended that you load the dataset into SPICE. However, if you want to use the `Select All` option with dataset parameters, this walkthrough describes a way to do so.

To begin, let's say you have a direct query dataset with custom SQL that contains a multivalue parameter called `States`:

```
select * from all_flights
```

```
where origin_state_abr in (<<$States>>)
```

To use the Select all option in a control that uses dataset parameters

1. In the **Parameters** pane of the analysis, find the dataset parameter that you want to use and choose **Edit** from the ellipsis (three dots) next to the parameter.
2. In the **Edit parameter** window that appears, enter a new default value in the **Static multiple default values** section. In our example, the default value is `All States`. Note that the example uses a leading space character so that the default value appears as the first item in the control.
3. Choose **Update** to update the parameter.
4. Navigate to the dataset that contains the dataset parameter that you're using in the analysis-by-analysis. Edit the custom SQL of the dataset to include a default use case for your new static multiple default values. Using the `All States` example, the SQL expression appears as follows:

```
select * from public.all_flights
where
  ' All States' in (<<$States>>) or
  origin_state_abr in (<<$States>>)
```

If the user chooses `All States` in the control, the new SQL expression returns all unique records. If the user chooses a different value from the control, the query returns values that were filtered by the dataset parameter.

Using controls with Select all and multivalue options

You can combine the previous `Select all` procedure with the multivalue control method discussed earlier to create dropdown controls that contain a `Select all` value in addition to multiple values that the user can select. This walkthrough assumes that you have followed the previous procedures, that you know how to map dataset parameters to analysis parameters, and that you can create controls in an analysis. For more information on mapping analysis parameters, see [Mapping dataset parameters in new Quick analyses](#). For more information on creating controls in an analysis that is using dataset parameters, see [Adding filter controls to mapped analysis parameters](#).

To add multiple values to a control with a Select all option and a mapped dataset parameter

1. Open the analysis that has the original dataset with a `Select all` custom SQL expression and a second dataset that includes all possible values of the filtered column that exists in the original dataset.
2. Navigate to the secondary dataset that was created earlier to return all values of a filtered column. Add a custom SQL expression that adds your previously configured `Select all` option to the query. The following example adds the `All States` record to the top of the list of returned values of the dataset:

```
(Select ' All States' as origin_state_abr)
Union All
(SELECT distinct origin_state_abr FROM all_flights
order by origin_state_abr asc)
```

3. Go back to the analysis that the datasets belong to and map the dataset parameter that you are using to the analysis parameter that you created in step 3 of the previous procedure. The analysis parameter and dataset parameter can have the same name. In our example, the analysis parameter is called `States`.
4. Create a new filter control or edit an existing filter control and choose **Hide Select All** to hide the disabled **Select All** option that appears in multivalue controls.

Once you create the control, users can use the same control to select all or multiple values of a filtered column in a dataset.

Using row-level security in Amazon Quick

Applies to: Enterprise Edition

In the Enterprise edition of Amazon Quick, you can restrict access to a dataset by configuring row-level security (RLS) on it. You can do this before or after you have shared the dataset. When you share a dataset with RLS with dataset owners, they can still see all the data. When you share it with readers, however, they can only see the data restricted by the permission dataset rules.

Also, when you embed Amazon Quick dashboards in your application for unregistered users of Quick, you can use row-level security (RLS) to filter/restrict data with tags. A tag is a user-specified string that identifies a session in your application. You can use tags to implement RLS controls for

your datasets. By configuring RLS-based restrictions in datasets, Quick filters the data based on the session tags tied to the user identity/session.

You can restrict access to a dataset using username or group-based rules, tag-based rules, or both.

Choose user-based rules if you want to secure data for users or groups provisioned (registered) in Quick. To do so, select a permissions dataset that contains rules set by columns for each user or group accessing the data. Only users or groups identified in the rules have access to data.

Choose tag-based rules only if you are using embedded dashboards and want to secure data for users not provisioned (unregistered users) in Quick. To do so, define tags on columns to secure data. Values to tags must be passed when embedding dashboards.

Topics

- [Using row-level security with user-based rules to restrict access to a dataset](#)
- [Using row-level security with tag-based rules to restrict access to a dataset when embedding dashboards for anonymous users](#)

Using row-level security with user-based rules to restrict access to a dataset

Applies to: Enterprise Edition

In the Enterprise edition of Amazon Quick, you can restrict access to a dataset by configuring row-level security (RLS) on it. You can do this before or after you have shared the dataset. When you share a dataset with RLS with dataset owners, they can still see all the data. When you share it with readers, however, they can only see the data restricted by the permission dataset rules. By adding row-level security, you can further control their access.

Note

When applying SPICE datasets to row-level security, each field in the dataset can contain up to 2,047 Unicode characters. Fields that contain more than this quota are truncated during ingestion. To learn more about SPICE data quotas, see [SPICE quotas for imported data](#).

To do this, you create a query or file with one column for user or group identification. You can use either `UserName` and `GroupName`, or alternatively `UserARN` and `GroupARN`. You can think of this as *adding a rule* for that user or group. Then you can add one column to the query or file for each field that you want to grant or restrict access to. For each user or group name that you add, you add the values for each field. You can use `NULL` (no value) to mean all values. To see examples of dataset rules, see [Creating dataset rules for row-level security](#).

To apply the dataset rules, you add the rules as a permissions dataset to your dataset. Keep in mind the following points:

- The permissions dataset can't contain duplicate values. Duplicates are ignored when evaluating how to apply the rules.
- Each user or group specified can see only the rows that *match* the field values in the dataset rules.
- If you add a rule for a user or group and leave all other columns with no value (`NULL`), you grant them access to all the data.
- If you don't add a rule for a user or group, that user or group can't see any of the data.
- The full set of rule records that are applied per user must not exceed 999. This limitation applies to the total number of rules that are directly assigned to a username, plus any rules that are assigned to the user through group names.
- If a field includes a comma (,) Amazon Quick treats each word separated from another by a comma as an individual value in the filter. For example, in ('AWS' , 'INC'), `AWS, INC` is considered as two strings: `AWS` and `INC`. To filter with `AWS, INC`, wrap the string with double quotation marks in the permissions dataset.

If the restricted dataset is a SPICE dataset, the number of filter values applied per user can't exceed 192,000 for each restricted field. This applies to the total number of filter values that are directly assigned to a username, plus any filter values that are assigned to the user through group names.

If the restricted dataset is a direct query dataset, the number of filter values applied per user varies from data sources.

Exceeding the filter value limit may cause the visual rendering to fail. We recommend adding an additional column to your restricted dataset to divide the rows into groups based on the original restricted column so that the filter list can be shortened.

Amazon Quick treats spaces as literal values. If you have a space in a field that you are restricting, the dataset rule applies to those rows. Amazon Quick treats both NULLs and blanks (empty strings "") as "no value". A NULL is an empty field value.

Depending on what data source your dataset is coming from, you can configure a direct query to access a table of permissions. Terms with spaces inside them don't need to be delimited with quotes. If you use a direct query, you can easily change the query in the original data source.

Or you can upload dataset rules from a text file or spreadsheet. If you are using a comma-separated value (CSV) file, don't include any spaces on the given line. Terms with spaces inside them need to be delimited with quotation marks. If you use dataset rules that are file-based, apply any changes by overwriting the existing rules in the dataset's permissions settings.

Datasets that are restricted are marked with the word **RESTRICTED** in the **Data** screen.

Child datasets that are created from a parent dataset that has RLS rules active retain the same RLS rules that the parent dataset has. You can add more RLS rules to the child dataset, but you can't remove the RLS rules that the dataset inherits from the parent dataset.

Child datasets that are created from a parent dataset that has RLS rules active can only be created with Direct Query. Child datasets that inherit the parent dataset's RLS rules aren't supported in SPICE.

Row-level security works only for fields containing textual data (string, char, varchar, and so on). It doesn't currently work for dates or numeric fields. Anomaly detection is not supported for datasets that use row-level security (RLS).

Creating dataset rules for row-level security

Use the following procedure to create a permissions file or query to use as dataset rules.

To create a permissions files or query to use as dataset rules

1. Create a file or a query that contains the dataset rules (permissions) for row-level security.

It doesn't matter what order the fields are in. However, all the fields are case-sensitive. Make sure that they exactly match the field names and values.

The structure should look similar to one of the following. Make sure that you have at least one field that identifies either users or groups. You can include both, but only one is required, and only one is used at a time. The field that you use for users or groups can have any name you choose.

Note

If you are specifying groups, use only Amazon Quick groups or Microsoft AD groups.

The following example shows a table with groups.

GroupName	Sales region	Segment	
EMEA-Sales	EMEA	Enterprise, SMB, Startup	
US-Sales	US	Enterprise	
US-Sales	US	SMB, Startup	
US-Sales	US	Startup	
APAC-Sales	APAC	Enterprise, SMB	
Corporate-Reporting			
APAC-Sales	APAC	Enterprise, Startup	

The following example shows a table with usernames.

UserName	Sales region	Segment	
AlejandroRosalez	EMEA	Enterprise, SMB, Startup	
MarthaRivera	US	Enterprise	
NikhilJayashankar	US	SMB, Startup	
PauloSantos	US	Startup	
SaanviSarkar	APAC	Enterprise, SMB	

UserName	Sales region	Segment	
sales-tps@example.com			
ZhangWei	APAC	Enterprise, Startup	

The following example shows a table with user and group Amazon Resource Names (ARNs).

UserARN	GroupARN	Sales region
arn:aws:quicksight:us-east-1:123456789012:user/default/Bob	arn:aws:quicksight:us-east-1:123456789012:group/default/group-1	APAC
arn:aws:quicksight:us-east-1:123456789012:user/default/Sam	arn:aws:quicksight:us-east-1:123456789012:group/default/group-2	US

Or if you use a .csv file, the structure should look similar to one of the following.

```

UserName,SalesRegion,Segment
AlejandroRosalez,EMEA,"Enterprise,SMB,Startup"
MarthaRivera,US,Enterprise
NikhilJayashankars,US,SMB
PauloSantos,US,Startup
SaanviSarkar,APAC,"SMB,Startup"
sales-tps@example.com,"",""
ZhangWei,APAC-Sales,"Enterprise,Startup"

```

```

GroupName,SalesRegion,Segment
EMEA-Sales,EMEA,"Enterprise,SMB,Startup"
US-Sales,US,Enterprise
US-Sales,US,SMB
US-Sales,US,Startup

```

```
APAC-Sales,APAC,"SMB,Startup"  
Corporate-Reporting,"",""  
APAC-Sales,APAC,"Enterprise,Startup"
```

```
UserARN,GroupARN,SalesRegion  
arn:aws:quicksight:us-east-1:123456789012:user/Bob,arn:aws:quicksight:us-  
east-1:123456789012:group/group-1,APAC  
arn:aws:quicksight:us-east-1:123456789012:user/Sam,arn:aws:quicksight:us-  
east-1:123456789012:group/group-2,US
```

Following is a SQL example.

```
/* for users*/  
select User as UserName, SalesRegion, Segment  
from tps-permissions;  
  
/* for groups*/  
select Group as GroupName, SalesRegion, Segment  
from tps-permissions;
```

2. Create a dataset for the dataset rules. To make sure that you can easily find it, give it a meaningful name, for example **Permissions-Sales-Pipeline**.

Rules Dataset flagging for row-level security

Use the following procedure to appropriately flag a dataset as a rules dataset.

Rules Dataset is a flag that distinguishes permission datasets used for row-level security from regular datasets. If a permissions dataset was applied to a regular dataset before March 31, 2025, it will have a Rules Dataset flag in the **Dataset** landing page.

If a permissions dataset was not applied to a regular dataset by March 31, 2025, it will be categorized as a regular dataset. To use it as a rules dataset, duplicate the permissions dataset and flag it as a rules dataset on the console when creating the dataset. Select EDIT DATASET and under the options, choose DUPLICATE AS RULES DATASET.

To successfully duplicate it as a rules dataset, ensure the original dataset has: 1. Required user metadata or group metadata column(s) and 2. Only string type columns.

To create a new rules dataset on the console, select NEW RULES DATASET under the NEW DATASET dropdown. When creating a rules dataset programmatically, add the following parameter: [UseAs:](#)

[RLS_RULES](#). This is an optional parameter that is only used to create a rules dataset. Once a dataset has been created, either through the console or programmatically, and flagged as either a rules dataset or a regular dataset, it cannot be changed.

Once datasets are flagged as rules datasets, Amazon Quick will apply strict SPICE ingestion rules on them. To ensure data integrity, SPICE ingestions for rules datasets will fail if there are invalid rows or cells exceeding length limits. You must fix the ingestion issues in order to re-initiate a successful ingestion. Strict ingestion rules are only applicable to rules datasets. Regular datasets will not have dataset ingestion failures when there are skipped rows or string truncations.

Applying row-level security

Use the following procedure to apply row-level security (RLS) by using a file or query as a dataset that contains the rules for permissions.

To apply row-level security by using a file or query

1. Confirm that you have added your rules as a new dataset. If you added them, but don't see them under the list of datasets, refresh the screen.
2. On the **Data** page, choose the dataset
3. On the dataset details page that opens, for **Row-level security**, choose **Set up**.
4. On the **Set up row-level security** page that opens, choose **User-based rules**.
5. From the list of datasets that appears, choose your permissions dataset.

If your permissions dataset doesn't appear on this screen, return to your datasets, and refresh the page.

6. For **Permissions policy** choose **Grant access to dataset**. Each dataset has only one active permissions dataset. If you try to add a second permissions dataset, it overwrites the existing one.

Important

Some restrictions apply to NULL and empty string values when working with row-level security:

- If your dataset has NULL values or empty strings ("") in the restricted fields, these rows are ignored when the restrictions are applied.

- Inside the permissions dataset, NULL values and empty strings are treated the same. For more information, see the following table.
- To prevent accidentally exposing sensitive information, Amazon Quick skips empty RLS rules that grant access to everyone. An *empty RLS rule* occurs when all columns of a row have no value. Quick RLS treats NULL, empty strings (""), or empty comma separated strings (for example ",",) as no value.
- After skipping empty rules, other nonempty RLS rules still apply.
- If a permission dataset has only empty rules and all of them were skipped, no one will have access to any data restricted by this permission dataset.

Rules for UserName, GroupName, SalesRegion, Segment	Granted access	
AlejandroRosalez,EMEA-Sales,EMEA,"Enterprise,SMB,Startup"	Sees all EMEA Enterprise, SMB, and Startup	
sales-tps@example.com,Corporate-Reporting,"", ""	Sees all rows	
User or group has no entry	Sees no rows	
","","",""	Skipped; sees no rows if all other rules are empty.	
NULL,"","",NULL	Skipped; sees no rows if all other rules are empty.	

Anyone whom you shared your dashboard with can see all the data in it, unless the dataset is restricted by dataset rules.

7. Choose **Apply dataset** to save your changes. Then, on the **Save data set rules?** page, choose **Apply and activate**. Changes in permissions apply immediately to existing users.
8. (Optional) To remove permissions, first remove the dataset rules from the dataset.

Make certain that the dataset rules are removed. Then, choose the permissions dataset and choose **Remove data set**.

To overwrite permissions, choose a new permissions dataset and apply it. You can reuse the same dataset name. However, make sure to apply the new permissions in the **Permissions** screen to make these permissions active. SQL queries dynamically update, so these can be managed outside of Amazon Quick. For queries, the permissions are updated when the direct query cache is automatically refreshed.

If you delete a file-based permissions dataset before you remove it from the target dataset, restricted users can't access the dataset. While the dataset is in this state, it remains marked as **RESTRICTED**. However, when you view **Permissions** for that dataset, you can see that it has no selected dataset rules.

To fix this, specify new dataset rules. Creating a dataset with the same name is not enough to fix this. You must choose the new permissions dataset on the **Permissions** screen. This restriction doesn't apply to direct SQL queries.

Using row-level security with tag-based rules to restrict access to a dataset when embedding dashboards for anonymous users

Applies to: Enterprise Edition

Intended audience: Amazon Quick Administrators and Amazon Quick developers

When you embed Amazon Quick dashboards in your application for users who are not provisioned (registered) in Quick, you can use row-level security (RLS) to filter/restrict data with tags. A tag is a user-specified string that identifies a session in your application. You can use tags to implement RLS controls for your datasets. By configuring RLS-based restrictions in datasets, Quick filters the data based on the session tags tied to the user identity/session.

For example, let's say you're a logistics company that has a customer-facing application for various retailers. Thousands of users from these retailers access your application to see metrics related to how their orders are getting shipped from your warehouse.

You don't want to manage thousands of users in Quick, so you use anonymous embedding to embed the selected dashboards in your application that your authenticated and authorized users can see. However, you want to make sure retailers see only data that is for their business and not for others. You can use RLS with tags to make sure your customers only see data that's relevant to them.

To do so, complete the following steps:

1. Add RLS tags to a dataset.
2. Assign values to those tags at runtime using the `GenerateEmbedUrlForAnonymousUser` API operation.

For more information about embedding dashboards for anonymous users using the `GenerateEmbedUrlForAnonymousUser` API operation, see [Embedding Amazon Quick Sight dashboards for anonymous \(unregistered\) users](#).

Before you can use RLS with tags, keep in mind the following points:

- Using RLS with tags is currently only supported for anonymous embedding, specifically for embedded dashboards that use the `GenerateEmbedUrlForAnonymousUser` API operation.
- Using RLS with tags isn't supported for embedded dashboards that use the `GenerateEmbedUrlForRegisteredUser` API operation or the old `GetDashboardEmbedUrl` API operation.
- RLS tags aren't supported with AWS Identity and Access Management (IAM) or the Quick identity type.
- When applying SPICE datasets to row-level security, each field in the dataset can contain up to 2,047 Unicode characters. Fields that contain more than this quota are truncated during ingestion. To learn more about SPICE data quotas, see [SPICE quotas for imported data](#).

Step 1: Add RLS tags to a dataset

You can add tag-based rules to a dataset in Amazon Quick. Alternatively, you can call the `CreateDataSet` or `UpdateDataSet` API operation and add tag-based rules that way. For more information, see [Add RLS tags to a dataset using the API](#).

Use the following procedure to add RLS tags to a dataset in Quick.

To add RLS tags to a dataset

1. From the Quick start page, choose **Data** at left.
2. Choose the dataset that you want to add RLS to.
3. On the dataset details page that opens, for **Row-level security**, choose **Set up**.
4. On the **Set up row-level security** page that opens, choose **Tag-based rules**.
5. For **Column**, choose a column that you want to add tag rules to.

For example, in the case for the logistics company, the `retailer_id` column is used.

Only columns with a string data type are listed.

6. For **Tag**, enter a tag key. You can enter any tag name that you want.

For example, in the case for the logistics company, the tag key `tag_retailer_id` is used. Doing this sets row-level security based on the retailer that's accessing the application.

7. (Optional) For **Delimiter**, choose a delimiter from the list, or enter your own.

You can use delimiters to separate text strings when assigning more than one value to a tag. The value for a delimiter can be 10 characters long, at most.

8. (Optional) For **Match all**, choose the `*`, or enter your own character or characters.

This option can be any character that you want to use when you want to filter by all the values in that column in the dataset. Instead of listing the values one by one, you can use the character. If this value is specified, it can be at least one character, or at most 256 characters long.

9. Choose **Add**.

The tag rule is added to the dataset and is listed at the bottom, but it isn't applied yet. To add another tag rule to the dataset, repeat steps 5–9. To edit a tag rule, choose the pencil icon that follows the rule. To delete a tag rule, choose the delete icon that follows the rule. You can add up to 50 tags to a dataset.

10. When you're ready to apply the tag rules to the dataset, choose **Apply rules**.
11. On the **Turn on tag-based security?** page that opens, choose **Apply and activate**.

The tag-based rules are now active. On the **Set up row-level security** page, a toggle appears for you to turn tag rules on and off for the dataset.

To turn off all tag-based rules for the dataset, switch the **Tag-Based rules** toggle off, and then enter "confirm" in the text box that appears.

On the **Data** page, a lock icon appears in the dataset row to indicate that tag rules are enabled.

You can now use tag rules to set tag values at runtime, described in [Step 2: Assign values to RLS tags at runtime](#). The rules only affect Quick readers when active.

Important

After tags are assigned and enabled on the dataset, make sure to give Quick authors permissions to see any of the data in the dataset when authoring a dashboard.

To give Quick authors permission to see data in the dataset, create a permissions file or query to use as dataset rules. For more information, see [Creating dataset rules for row-level security](#).

After you create a tag-based rule, a new **Manage rules** table appears that shows how your tag-based rules relate to each other. To make changes to the rules listed in the **Manage rules** table, choose the pencil icon that follows the rule. Then add or remove tags, and choose **Update**. To apply your updated rule to the dataset, choose **Apply**.

(Optional) Add the OR condition to RLS tags

You can also add the OR condition to your tag-based rules to further customize the way data is presented to your Quick account users. When you use the OR condition with your tag-based rules, visuals in Quick appear if at least one tag defined in the rule is valid.

To add the OR condition to your tag-based rules

1. In the **Manage rules** table, choose **Add OR condition**.
2. In the **Select tag** dropdown list that appears, choose the tag that you want to create an OR condition for. You can add up to 50 OR conditions to the **Manage rules** table. You can add multiple tags to a single column in a dataset, but at least one column tag needs to be included in a rule.
3. Choose **Update** to add the condition to your rule, then choose **Apply** to apply the updated rule to your dataset.

Add RLS tags to a dataset using the API

Alternatively, you can configure and enable tag-based row-level security on your dataset by calling the `CreateDataSet` or `UpdateDataSet` API operation. Use the following examples to learn how.

Important

When configuring session tags in the API call,

- Treat session tags as security credentials. Do not expose session tags to end users or client-side code.
- Implement server-side controls. Ensure that session tags are set exclusively by your trusted backend services, not by parameters that end users can modify.
- Protect session tags from enumeration. Ensure that users in one tenant cannot discover or guess sessionTag values belonging to other tenants.
- Review your architecture. If downstream customers or partners are allowed to call the API directly, evaluate whether those parties could specify sessionTag values for tenants they should not access.

CreateDataSet

The following is an example for creating a dataset that uses RLS with tags. It assumes the scenario of the logistics company described previously. The tags are defined in the `row-level-permission-tag-configuration` element. The tags are defined on the columns that you want to secure the data for. For more information about this optional element, see [RowLevelPermissionTagConfiguration](#) in the *Amazon Quick API Reference*.

```
create-data-set
  --aws-account-id <value>
  --data-set-id <value>
  --name <value>
  --physical-table-map <value>
  [--logical-table-map <value>]
  --import-mode <value>
  [--column-groups <value>]
  [--field-folders <value>]
  [--permissions <value>]
  [--row-level-permission-data-set <value>]
  [--column-level-permission-rules <value>]
```

```

[--tags <value>]
[--cli-input-json <value>]
[--generate-cli-skeleton <value>]
[--row-level-permission-tag-configuration
'{
  "Status": "ENABLED",
  "TagRules":
  [
    {
      "TagKey": "tag_retailer_id",
      "ColumnName": "retailer_id",
      "TagMultiValueDelimiter": ",",
      "MatchAllValue": "*"
    },
    {
      "TagKey": "tag_role",
      "ColumnName": "role"
    }
  ],
  "TagRuleConfigurations":
  [
    tag_retailer_id
  ],
  [
    tag_role
  ]
}']
]

```

The tags in this example are defined in the TagRules part of the element. In this example, two tags are defined based on two columns:

- The tag_retailer_id tag key is defined for the retailer_id column. In this case for the logistics company, this sets row-level security based on the retailer that's accessing the application.
- The tag_role tag key is defined for the role column. In this case for the logistics company, this sets an additional layer of row-level security based on the role of the user accessing your application from a specific retailer. An example is store_supervisor or manager.

For each tag, you can define TagMultiValueDelimiter and MatchAllValue. These are optional.

- **TagMultiValueDelimiter** – This option can be any string that you want to use to delimit the values when you pass them at runtime. The value can be 10 characters long, at most. In this case, a comma is used as the delimiter value.
- **MatchAllValue** – This option can be any character that you want to use when you want to filter by all the values in that column in the dataset. Instead of listing the values one by one, you can use the character. If specified, this value can be at least one character, or at most 256 characters long. In this case, an asterisk is used as the match all value.

While configuring the tags for dataset columns, turn them on or off using the mandatory property **Status**. For enabling the tag rules use the value **ENABLED** for this property. By turning on tag rules, you can use them to set tag values at runtime, described in [Step 2: Assign values to RLS tags at runtime](#).

The following is an example of the response definition.

```
{
  "Status": 201,
  "Arn": "arn:aws:quicksight:us-west-2:11112222333:dataset/RLS-Dataset",
  "DataSetId": "RLS-Dataset",
  "RequestId": "aa4f3c00-b937-4175-859a-543f250f8bb2"
}
```

UpdateDataSet

UpdateDataSet

You can use the `UpdateDataSet` API operation to add or update RLS tags for an existing dataset.

The following is an example of updating a dataset with RLS tags. It assumes the scenario of the logistics company described previously.

```
update-data-set
--aws-account-id <value>
--data-set-id <value>
--name <value>
--physical-table-map <value>
[--logical-table-map <value>]
--import-mode <value>
```

```

[--column-groups <value>
[--field-folders <value>]
[--row-level-permission-data-set <value>]
[--column-level-permission-rules <value>]
[--cli-input-json <value>]
[--generate-cli-skeleton <value>]
[--row-level-permission-tag-configuration
'{
  "Status": "ENABLED",
  "TagRules":
  [
    {
      "TagKey": "tag_retailer_id",
      "ColumnName": "retailer_id",
      "TagMultiValueDelimiter": ",",
      "MatchAllValue": "*"
    },
    {
      "TagKey": "tag_role",
      "ColumnName": "role"
    }
  ],
  "TagRuleConfigurations":
  [
    tag_retailer_id
  ],
  [
    tag_role
  ]
}']
]
```

The following is an example of the response definition.

```

{
  "Status": 201,
  "Arn": "arn:aws:quicksight:us-west-2:11112222333:dataset/RLS-Dataset",
  "DataSetId": "RLS-Dataset",
  "RequestId": "aa4f3c00-b937-4175-859a-543f250f8bb2"
}
```


⚠ Important

After tags are assigned and enabled on the dataset, make sure to give Quick authors permissions to see any of the data in the dataset when authoring a dashboard.

To give Quick authors permission to see data in the dataset, create a permissions file or query to use as dataset rules. For more information, see [Creating dataset rules for row-level security](#).

For more information about the `RowLevelPermissionTagConfiguration` element, see [RowLevelPermissionTagConfiguration](#) in the *Amazon Quick API Reference*.

Step 2: Assign values to RLS tags at runtime

You can use tags for RLS only for anonymous embedding. You can set values for tags using the `GenerateEmbedUrlForAnonymousUser` API operation.

⚠ Important

When configuring session tags in the API call,

- Treat session tags as security credentials. Do not expose session tags to end users or client-side code.
- Implement server-side controls. Ensure that session tags are set exclusively by your trusted backend services, not by parameters that end users can modify.
- Protect session tags from enumeration. Ensure that users in one tenant cannot discover or guess sessionTag values belonging to other tenants.
- Review your architecture. If downstream customers or partners are allowed to call the API directly, evaluate whether those parties could specify sessionTag values for tenants they should not access.

The following example shows how to assign values to RLS tags that were defined in the dataset in the previous step.

```
POST /accounts/AwsAccountId/embed-url/anonymous-user
HTTP/1.1
Content-type: application/json
{
```

```

"AwsAccountId": "string",
"SessionLifetimeInMinutes": integer,
"Namespace": "string", // The namespace to which the anonymous end user virtually
belongs
"SessionTags": // Optional: Can be used for row-level security
[
  {
    "Key": "tag_retailer_id",
    "Value": "West,Central,South"
  }
  {
    "Key": "tag_role",
    "Value": "shift_manager"
  }
],
"AuthorizedResourceArns":
[
  "string"
],
"ExperienceConfiguration":
{
  "Dashboard":
  {
    "InitialDashboardId": "string"
    // This is the initial dashboard ID the customer wants the user to land on. This
ID goes in the output URL.
  }
}
}

```

The following is an example of the response definition.

```

HTTP/1.1 Status
Content-type: application/json

{
  "EmbedUrl": "string",
  "RequestId": "string"
}

```

RLS support without registering users in Quick is supported only in the `GenerateEmbedUrlForAnonymousUser` API operation. In this operation, under `SessionTags`, you can define the values for the tags associated with the dataset columns.

In this case, the following assignments are defined:

- Values West, Central, and South are assigned to the `tag_retailer_id` tag at runtime. A comma is used for the delimiter, which was defined in `TagMultipleValueDelimiter` in the dataset. To use all values in the column, you can set the value to `*`, which was defined as the `MatchAllValue` when creating the tag.
- The value `shift_manager` is assigned to the `tag_role` tag.

The user using the generated URL can view only the rows having the `shift_manager` value in the `role` column. That user can view only the value West, Central, or South in the `retailer_id` column.

For more information about embedding dashboards for anonymous users using the `GenerateEmbedUrlForAnonymousUser` API operation, see [Embedding Amazon Quick Sight dashboards for anonymous \(unregistered\) users](#), or [GenerateEmbedUrlForAnonymousUser](#) in the *Amazon Quick API Reference*

Using column-level security to restrict access to a dataset

In the Enterprise edition of Quick, you can restrict access to a dataset by configuring column-level security (CLS) on it. A dataset or analysis with CLS enabled has the restricted



symbol next to it.

By default, all users and groups have access to the data. By using CLS, you can manage access to specific columns in your dataset.

If you use an analysis or dashboard that contains datasets with CLS restrictions that you don't have access to, you can't create, view, or edit visuals that use the restricted fields. For most visual types, if a visual has restricted columns that you don't have access to, you can't see the visual in your analysis or dashboard.

Tables and pivot tables behave differently. If a table or pivot table uses restricted columns in the **Rows** or **Columns** field wells, and you don't have access to these restricted columns, you can't see the visual in an analysis or dashboard. If a table or pivot table has restricted columns in the **Values** field well, you can see the table in an analysis or dashboard with only the values that you have access to. The values for restricted columns show as Not Authorized.

To enable column-level security on an analysis or dashboard, you need administrator access.

To create a new analysis with CLS

1. On the Quick start page, choose the **Analyses** tab.
2. At upper right, choose **New analysis**.
3. Choose a dataset, and choose **Column-level security**.
4. Select the columns that you want to restrict, and then choose **Next**. By default, all groups and users have access to all columns.
5. Choose who can access each column, and then choose **Apply** to save your changes.

To use an existing analysis for CLS

1. On the Quick start page, choose the **Data** tab.
2. On the Data page, open your dataset
3. On the dataset details page that opens, for **Column-level security**, choose **Set up**.
4. Select the columns that you want to restrict, and then choose **Next**. By default, all groups and users have access to all columns.
5. Choose who can access each column, and then choose **Apply** to save your changes.

To create a dashboard with CLS

1. On the Quick navigation pane, choose the **Analyses** tab.
2. Choose the analysis that you want to create a dashboard of.
3. At upper right, choose **Publish**.
4. Choose one of the following:
 - To create a new dashboard, choose **Publish new dashboard as** and enter a name for the new dashboard.
 - To replace an existing dashboard, choose **Replace an existing dashboard** and choose the dashboard from the list.

Additionally, you can choose **Advanced publish options**. For more information, see [Publishing dashboards](#).

5. Choose **Publish dashboard**.
6. (Optional) Do one of the following:

- To publish a dashboard without sharing, choose **x** at the upper right of the **Share dashboard with users** screen when it appears. You can share the dashboard later by choosing **Share** from the application bar.
- To share the dashboard, follow the procedure in [Sharing Amazon Quick Sight dashboards](#).

Running queries as an IAM role in Amazon Quick

You can enhance data security by using fine-grained access policies rather than broader permissions for data sources connected to Amazon Athena, Amazon Redshift or Amazon S3. You start by creating an AWS Identity and Access Management (IAM) role with permissions to be activated when a person or an API starts a query. Then, an Quick administrator or a developer assigns the IAM Role to an Athena or Amazon S3 data source. With the role in place, any person or API that runs the query has the exact permissions necessary to run the query.

Here are some things to consider before you commit to implementing run-as roles to enhance data security:

- Articulate how the additional security works to your advantage.
- Work with your Quick administrator to learn if adding roles to data sources helps you to better meet your security goals or requirements.
- Ask if this type of security, for the number of data sources and people and applications involved, can be feasibly documented and maintained by your team? If not, then who will undertake that part of the work?
- In a structured organization, locate stakeholders in parallel teams in Operations, Development, and IT Support. Ask for their experience, advice, and willingness to support your plan.
- Before you launch your project, consider doing a proof of concept that involves the people who need access to the data.

The following rules apply to using run-as roles with Athena, Amazon Redshift, and Amazon S3:

- Each data source can have only one associated RoleArn. Consumers of the data source, who typically access datasets and visuals, can generate many different types of queries. The role places boundaries on which queries work and which don't work.
- The ARN must correspond to an IAM role in the same AWS account as the Quick instance that uses it.

- The IAM role must have a trust relationship allowing Quick to assume the role.
- The identity that calls Quick's APIs must have permission to pass the role before they can update the `RoleArn` property. You only need to pass the role when creating or updating the role ARN. The permissions aren't re-evaluated later on. Similarly, the permission isn't required when the role ARN is omitted.
- When the role ARN is omitted, the Athena or Amazon S3 data source uses the account-wide role and scope-down policies.
- When the role ARN is present, the account-wide role and any scope-down policies are both ignored. For Athena data sources, Lake Formation permissions are not ignored.
- For Amazon S3 data sources, both the manifest file and the data specified by the manifest file must be accessible using the IAM role.
- The ARN string needs to match an existing IAM role in the AWS account and AWS Region where the data is located and queried.

When Quick connects to another service in AWS, it uses an IAM role. By default, this less granular version of the role is created by Quick for each service it uses, and the role is managed by AWS account administrators. When you add an IAM role ARN with a custom permissions policy, you override the broader role for your data sources that need extra protection. For more information about policies, see [Create a customer managed policy](#) in the IAM User Guide.

Run queries with Athena data sources

Use the API to attach the ARN to the Athena data source. To do so, add the role ARN in the [RoleArn](#) property of [AthenaParameters](#). For verification, you can see the role ARN on the **Edit Athena data source** dialog box. However, **Role ARN** is a read-only field.

To get started, you need a custom IAM role, which we demonstrate in the following example.

Keep in mind that the following code example is for learning purposes only. Use this example in a temporary development and testing environment only, and not in a production environment. The policy in this example doesn't secure any specific resource, which must be in a deployable policy. Also, even for development, you need to add your own AWS account information.

The following commands create a simple new role and attach a few policies that grant permissions to Quick.

```
aws iam create-role \  
    --role-name TestAthenaRoleForQuickSight \  
    --policy-names QuickSightAthenaPolicy,QuickSightS3Policy
```

```
--description "Test Athena Role For QuickSight" \
--assume-role-policy-document '{
  "Version": "2012-10-17"
  ,
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": "quicksight.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}'
```

After you've identified or created an IAM role to use with each data source, attach the policies by using the `attach-role-policy`.

```
aws iam attach-role-policy \
  --role-name TestAthenaRoleForQuickSight \
  --policy-arn arn:aws:iam::222222222222:policy/service-role/
AWSQuickSightS3Policy1

aws iam attach-role-policy \
  --role-name TestAthenaRoleForQuickSight \
  --policy-arn arn:aws:iam::aws:policy/service-role/AWSQuicksightAthenaAccess1

aws iam attach-role-policy \
  --role-name TestAthenaRoleForQuickSight \
  --policy-arn arn:aws:iam::aws:policy/AmazonS3Access1
```

After you verify your permissions, you can use the role in Quick data sources by creating a new role or updating an existing role. When using these commands, update the AWS account ID and AWS Region to match your own.

Remember, these example code snippets are not for production environments. AWS strongly recommends that you identify and use a set of least privilege policies for your production cases.

```
aws quicksight create-data-source
  --aws-account-id 222222222222 \
  --region us-east-1 \
  --data-source-id "athena-with-custom-role" \
```

```
--cli-input-json '{
  "Name": "Athena with a custom Role",
  "Type": "ATHENA",
  "data sourceParameters": {
    "AthenaParameters": {
      "RoleArn": "arn:aws:iam::222222222222:role/
TestAthenaRoleForQuickSight"
    }
  }
}'
```

Run queries with Amazon Redshift data sources

Connect your Amazon Redshift data with the run-as role to enhance your data security with fine-grained access policies. You can create a run-as role for Amazon Redshift data sources that use a public network or a VPC connection. You specify the connection type that you want to use in the **Edit Amazon Redshift data source** dialog box. The run-as role is not supported for Amazon Redshift Serverless data sources.

To get started, you need a custom IAM role, which we demonstrate in the following example. The following commands create a sample new role and attach policies that grant permissions to Quick.

```
aws iam create-role \
--role-name TestRedshiftRoleForQuickSight \
--description "Test Redshift Role For QuickSight" \
--assume-role-policy-document '{
  "Version": "2012-10-17"
  ,
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": "quicksight.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}'
```

After you identify or create an IAM role to use with each data source, attach the policies with an `attach-role-policy`. If the `redshift:GetClusterCredentialsWithIAM` permission is attached to the role that you want to use, the values for `DatabaseUser` and `DatabaseGroups` are optional.


```
aws iam attach-role-policy \
--role-name TestRedshiftRoleForQuickSight \
--policy-arn arn:aws:iam:111122223333:policy/service-role/AWSQuickSightRedshiftPolicy

aws iam create-policy --policy-name RedshiftGetClusterCredentialsPolicy1 \
--policy-document file://redshift-get-cluster-credentials-policy.json

aws iam attach-role-policy \
--role-name TestRedshiftRoleForQuickSight \
--policy-arn arn:aws:iam:111122223333:policy/RedshiftGetClusterCredentialsPolicy1
// redshift-get-cluster-credentials-policy.json
{
  "Version": "2012-10-17"      ,
  "Statement": [
    {
      "Sid": "RedshiftGetClusterCredentialsPolicy",
      "Effect": "Allow",
      "Action": [
        "redshift:GetClusterCredentials"
      ],
      "Resource": [
        "*"
      ]
    }
  ]
}
```

The example above creates a data source that uses the RoleARN, DatabaseUser, and DatabaseGroups IAM parameters. If you want to establish the connection only through the IAM RoleARN parameter, attach the `redshift:GetClusterCredentialsWithIAM` permission to your role, shown in the example below.

```
aws iam attach-role-policy \
--role-name TestRedshiftRoleForQuickSight \
--policy-arn arn:aws:iam:111122223333:policy/RedshiftGetClusterCredentialsPolicy1 //
redshift-get-cluster-credentials-policy.json {
  "Version": "2012-10-17"      ,
  "Statement": [
    {
      "Sid": "RedshiftGetClusterCredentialsPolicy",
```

```

        "Effect": "Allow",
        "Action": [ "redshift:GetClusterCredentialsWithIAM" ],
        "Resource": [ "*" ]
    }
]
}"

```

After you verify your permissions, you can use the role in Quick data sources by creating a new role or updating an existing role. When using these commands, update the AWS account ID and AWS Region to match your own.

```

aws quicksight create-data-source \
--region us-west-2 \
--endpoint https://quicksight.us-west-2.quicksight.aws.com/ \
--cli-input-json file://redshift-data-source-iam.json \
redshift-data-source-iam.json is shown as below
{
  "AwsAccountId": "AWSACCOUNTID",
  "DataSourceId": "DATSOURCEID",
  "Name": "Test redshift demo iam",
  "Type": "REDSHIFT",
  "DataSourceParameters": {
    "RedshiftParameters": {
      "Database": "integ",
      "Host": "redshiftdemocluster.us-west-2.redshift.amazonaws.com",
      "Port": 8192,
      "ClusterId": "redshiftdemocluster",
      "IAMParameters": {
        "RoleArn":
"arn:aws:iam::222222222222:role/TestRedshiftRoleForQuickSight",
        "DatabaseUser": "user",
        "DatabaseGroups": ["admin_group", "guest_group", "guest_group_1"]
      }
    }
  },
  "Permissions": [
    {
      "Principal": "arn:aws:quicksight:us-east-1:AWSACCOUNTID:user/default/demoname",
      "Actions": [
        "quicksight:DescribeDataSource",
        "quicksight:DescribeDataSourcePermissions",
        "quicksight:PassDataSource",
        "quicksight:UpdateDataSource",

```

```

        "quicksight:DeleteDataSource",
        "quicksight:UpdateDataSourcePermissions"
    ]
}
]
}

```

If your data source uses the VPC connection type, use the following VPC configuration.

```

{
  "AwsAccountId": "AWSACCOUNTID",
  "DataSourceId": "DATSOURCEID",
  "Name": "Test redshift demo iam vpc",
  "Type": "REDSHIFT",
  "DataSourceParameters": {
    "RedshiftParameters": {
      "Database": "mydb",
      "Host": "vpcdemo.us-west-2.redshift.amazonaws.com",
      "Port": 8192,
      "ClusterId": "vpcdemo",
      "IAMParameters": {
        "RoleArn":
"arn:aws:iam::222222222222:role/TestRedshiftRoleForQuickSight",
        "DatabaseUser": "user",
        "AutoCreateDatabaseUser": true
      }
    }
  },
  "VpcConnectionProperties": {
    "VpcConnectionArn": "arn:aws:quicksight:us-west-2:222222222222:vpcConnection/VPC
Name"
  },
  "Permissions": [
    {
      "Principal": "arn:aws:quicksight:us-east-1:222222222222:user/default/demoname",
      "Actions": [
        "quicksight:DescribeDataSource",
        "quicksight:DescribeDataSourcePermissions",
        "quicksight:PassDataSource",
        "quicksight:UpdateDataSource",
        "quicksight:DeleteDataSource",
        "quicksight:UpdateDataSourcePermissions"
      ]
    }
  ]
}

```

```
}  
]  
}
```

If your data source uses the `redshift:GetClusterCredentialsWithIAM` permission and doesn't use the `DatabaseUser` or `DatabaseGroups` parameters, grant the role access to some or all tables in the schema. To see if a role has been granted `SELECT` permissions to a specific table, input the following command into the Amazon Redshift Query Editor.

```
SELECT  
u.username,  
t.schemaname||'.'||t.tablename,  
has_table_privilege(u.username,t.tablename,'select') AS user_has_select_permission  
FROM  
pg_user u  
CROSS JOIN  
pg_tables t  
WHERE  
u.username = 'IAMR:RoLeName'  
AND t.tablename = tableName
```

For more information about the `SELECT` action in the Amazon Redshift Query Editor, see [SELECT](#).

To grant `SELECT` permissions to the role, input the following command in the Amazon Redshift Query Editor.

```
GRANT SELECT ON { [ TABLE ] table_name [, ...] | ALL TABLES IN SCHEMA  
schema_name [, ...] } TO "IAMR:RoLeName";
```

For more information about the `GRANT` action in the Amazon Redshift Query Editor, see [GRANT](#).

Run queries with Amazon S3 data sources

Amazon S3 data sources contain a manifest file that Quick uses to find and parse your data. You can upload a JSON manifest file through the Quick console, or you can provide a URL that points to a JSON file in an S3 bucket. If you choose to provide a URL, Quick must be granted permission to access the file in Amazon S3. Use the Quick administration console to control access to the manifest file and the data that it references.

With the **RoleArn** property, you can grant access to the manifest file and the data that it references through a custom IAM role that overrides the account-wide role. Use the API to attach the ARN to the manifest file of the Amazon S3 data source. To do so, include the role ARN in the [RoleArn](#) property of [S3Parameters](#). For verification, you can see the role ARN in the **Edit S3 data source** dialog box. However, **Role ARN** is a read-only field, as shown in the following screenshot.

To get started, create an Amazon S3 manifest file. Then, you can either upload it to Amazon Quick when you create a new Amazon S3 dataset or place the file into the Amazon S3 bucket that contains your data files. View the following example to see what a manifest file might look like:

```
{
  "fileLocations": [
    {
      "URIPrefixes": [
        "s3://quicksightUser-run-as-role/data/"
      ]
    }
  ],
  "globalUploadSettings": {
    "format": "CSV",
    "delimiter": ",",
    "textqualifier": "'",
    "containsHeader": "true"
  }
}
```

For instructions on how to create a manifest file, see [Supported formats for Amazon S3 manifest files](#).

After you have created a manifest file and added it to your Amazon S3 bucket or uploaded it to Quick, create or update an existing role in IAM that grants `s3:GetObject` access. The following example illustrates how to update an existing IAM role with the AWS API:

```
aws iam put-role-policy \
  --role-name QuickSightAccessToS3RunAsRoleBucket \
  --policy-name GrantS3RunAsRoleAccess \
  --policy-document '{
    "Version": "2012-10-17"
    ,
    "Statement": [
      {
        "Effect": "Allow",
```

```

        "Action": "s3:ListBucket",
        "Resource": "arn:aws:s3:::s3-bucket-name"
    },
    {
        "Effect": "Allow",
        "Action": "s3:GetObject",
        "Resource": "arn:aws:s3:::s3-bucket-name/manifest.json"
    },
    {
        "Effect": "Allow",
        "Action": "s3:GetObject",
        "Resource": "arn:aws:s3:::s3-bucket-name/*"
    }
]
}'

```

After your policy grants `s3:GetObject` access, you can begin creating data sources that apply the updated `put-role-policy` to the Amazon S3 data source's manifest file.

```

aws quicksight create-data-source --aws-account-id 111222333444 --region us-west-2 --
endpoint https://quicksight.us-west-2.quicksight.aws.com/ \
--data-source-id "s3-run-as-role-demo-source" \
--cli-input-json '{
    "Name": "S3 with a custom Role",
    "Type": "S3",
    "DataSourceParameters": {
        "S3Parameters": {
            "RoleArn": "arn:aws:iam::111222333444:role/
QuickSightAccessRunAsRoleBucket",
            "ManifestFileLocation": {
                "Bucket": "s3-bucket-name",
                "Key": "manifest.json"
            }
        }
    }
}'

```

After you verify your permissions, you can use the role in Quick data sources, either by creating a new role or updating an existing role. When using these commands, be sure to update the AWS account ID and AWS Region to match your own.

Deleting datasets

Important

Currently, deleting a dataset is irreversible and can cause irreversible loss of work. Deletes don't cascade to delete dependent objects. Instead, dependent objects stop working, even if you replace the deleted dataset with an identical dataset.

Before you delete a dataset, we strongly recommend that you first point each dependent analysis or dashboard to a new dataset.

Currently, when you delete a dataset while dependent visuals still exist, the analyses and dashboards that contain those visuals have no way to assimilate new metadata. They remain visible, but they can't function. They can't be repaired by adding an identical dataset.

This is because datasets include metadata that is integral to the analyses and dashboards that depend on that dataset. This metadata is uniquely generated for each dataset. Although the Quick Sight engine can read the metadata, it isn't readable by humans (for example, it doesn't contain field names). So, an exact replica of the dataset has different metadata. Each dataset's metadata is unique, even for multiple datasets that share the same name and the same fields.

To delete a dataset

1. Make sure that the dataset isn't being used by any analysis or dashboard that someone wants to keep using.

On the **Data** page, choose the dataset that you no longer need. Then choose **Delete Dataset** at upper-right.

2. If you receive a warning if this dataset is in use, track down all dependent analyses and dashboards and point them at a different dataset. If this isn't feasible, try one or more of these best practices instead of deleting it:
 - Rename the dataset, so that the dataset is clearly deprecated.
 - Filter the data, so that the dataset has no rows.
 - Remove everyone else's access to the dataset.

We recommend that you use whatever means you can to inform owners of dependent objects that this dataset is being deprecated. Also, make sure that you provide sufficient time for them to take action.

3. After you make sure that there are no dependent objects that will stop functioning after the dataset is deleted, choose the dataset and choose **Delete Data Set**. Confirm your choice, or choose **Cancel**.

Important

Currently, deleting a dataset is irreversible and can cause irreversible loss of work. Deletes don't cascade to delete dependent objects. Instead, dependent objects stop working, even if you replace the deleted dataset with an identical dataset.

Adding a dataset to an analysis

After you have created an analysis, you can add more datasets to the analysis. Then, you can use them to create more visuals.

From within the analysis, you can open any dataset for editing, for example to add or remove fields, or perform other data preparation. You can also remove or replace data sets.

The currently selected dataset displays at the top of the **Data** pane. This is the dataset that is used by the currently selected visual. Each visual can use only one dataset. Choosing a different visual changes the selected dataset to the one used by that visual.

To change the selected dataset manually, choose the dataset list at the top of the **Data** pane and then choose a different dataset. This deselects the currently selected visual if it doesn't use this dataset. Then, choose a visual that uses the selected dataset. Or choose **Add** in the **Visuals** pane to create a new visual using the selected dataset.

If you choose **Suggested** on the tool bar to see suggested visuals, you'll see visuals based on the currently selected dataset.

Only filters for the currently selected dataset are shown in the **Filter** pane, and you can only create filters on the currently selected dataset.

Topics

- [Replacing datasets](#)
- [Remove a dataset from an analysis](#)

Use the following procedure to add a dataset to an analysis or edit a dataset used by an analysis.

To add a dataset to an analysis

1. On the analysis page, navigate to the **Data** pane and expand the **Dataset** dropdown.
2. Choose **Add a new dataset** to add a dataset. Or, choose **Manage datasets** to edit a dataset. For more information about editing a dataset, see [Editing datasets](#).
3. A list of your datasets appears. Choose a dataset and then choose **Select**. To cancel, choose **Cancel**.

Replacing datasets

In an analysis, you can add, edit, replace, or remove datasets. Use this section to learn how to replace your dataset.

When you replace a dataset, the new dataset should have similar columns, if you expect the visual to work the way you designed it. Replacing the dataset also clears the undo and redo history for the analysis. This means you can't use the undo and redo buttons on the application bar to navigate your changes. So, when you decide to change the dataset, your analysis design should be somewhat stable—not in the middle of an editing phase.

To replace a dataset

1. On the analysis page, navigate to the **Data** pane and expand the **Dataset** dropdown.
2. Choose **Manage datasets**.
3. Choose the ellipsis (three dots) next to the dataset that you want to replace, and then choose **Replace**.
4. In the **Select replacement dataset** page, choose a dataset from the list, and then choose **Select**.

Note

Replacing a dataset clears the undo and redo history for this analysis.

The dataset is replaced with the new one. The field list and visuals are updated with the new dataset.

At this point, you can choose to add a new dataset, edit the new dataset, or replace it with a different one. Choose **Close** to exit.

If your new dataset doesn't match

In some cases, the selected replacement dataset doesn't contain all of the fields and hierarchies used by the visuals, filters, parameters, and calculated fields in your analysis. If so, you receive a warning from Quick Sight that shows a list of mismatched or missing columns.

If this happens, you can update the field mapping between the two datasets.

To update the field mapping

1. In the **Mismatch in replacement dataset** page, choose **Update field mapping**.
2. In the **Update field mapping** page, choose the drop-down menu for the field(s) you want to map and choose a field from the list to map it to.

If the field is missing from the new dataset, choose **Ignore this field**.

3. Choose **Confirm** to confirm your updates.
4. Choose **Close** to close the page and return to your analysis.

The dataset is replaced with the new one. The fields list and visuals are updated with the new dataset.

Any visuals that were using a field that's now missing from the new dataset update to blank. You can readd fields to the visual or remove the visual from your analysis.

If you change your mind after replacing the dataset, you can still recover. Let's say you replace the dataset and then find that it's too difficult to change your analysis to match the new dataset. You can undo any changes you made to your analysis. You can then replace the new dataset with the original one, or with a dataset that more closely matches the requirements of the analysis.

Remove a dataset from an analysis

Use the following procedure to delete a dataset from an analysis.

To delete a dataset from an analysis

1. On the analysis page, navigate to the **Data** pane and expand the **Dataset** dropdown.
2. Choose **Manage datasets**.
3. Choose the ellipsis (three dots) next to the dataset that you want to replace, and then choose **Remove**. You can't delete a dataset if it's the only one in the analysis.
4. Choose **Close** to close the dialog box.

Working with data sources in Amazon Quick Sight

Use a data source to access an external data store. Amazon S3 data sources save the manifest file information. In contrast, Salesforce and database data sources save connection information like credentials. In such cases, you can easily create multiple datasets from the data store without having to re-enter information. Connection information isn't saved for text or Microsoft Excel files.

Topics

- [Creating a data source](#)
- [Editing a data source](#)
- [Deleting a data source](#)

Creating a data source

Intended audience: Amazon Quick authors

As an analysis author in Amazon Quick, you don't need to understand anything about the infrastructure that you use to connect to your data. You set up a new data source only once.

After a data source is set up, you can access it from its tile in the Quick console. You can use it to create one or more datasets. After a dataset is set up, you can also access the dataset from its tile. By abstracting away the technical details, Amazon Quick Sight simplifies data connections.

Note

You don't need to store connection settings for files that you plan to upload manually. For more information about file uploads, see [Creating datasets](#).

Before you begin adding a new data-source connection profile to Amazon Quick, first collect the information that you need to connect to the data source. In some cases, you might plan to copy and paste settings from a file. If so, make sure that the file doesn't contain formatting characters (list bullets or numbers) or blank space characters (spaces, tabs). Also make sure that the file doesn't contain nontext "gremlin" characters such as non-ASCII, null (ASCII 0), and control characters.

The following list includes the information to collect the most commonly used settings:

- The data source to connect to.

Make sure that you know which source that you need to connect to for reporting. This source might be different than the source that stores, processes, or provides access to the data.

For example, let's say that you're a new analyst in a large company. You want to analyze data from your ordering system, which you know uses Oracle. However, you can't directly query the online transaction processing (OLTP) data. A subset of data is extracted and stored in a bucket on Amazon S3, but you don't have access to that either. Your new co-workers explain that they use AWS Glue crawlers to read the files and AWS Lake Formation to access them. With more research, you learn that you need to use an Amazon Athena query as your data source in Amazon Quick Sight. The point here is that it isn't always obvious which type of data source to choose.

- A descriptive name for the new data source tile.

Each new data source connection needs a unique and descriptive name. This name displays on the Amazon Quick Sight list of existing data sources, which is at the bottom of the **Create a Data Set** screen. Use a name that makes it easy to distinguish your data sources from other similar data sources. Your new Amazon Quick Sight data source profile displays both the database software logo and the custom name that you assign.

- The name of the server or instance to connect to.

A unique name or other identifier identifies the server connector of the data source on your network. The descriptors vary depending on which one you're connecting to, but it's usually one or more of the following:

- Hostname
- IP address
- Cluster ID
- Instance ID

- Connector
- Site-based URL
- The name of the collection of data that you want to use.

The descriptor varies depending on the data source, but it's usually one of the following:

- Database
- Warehouse
- S3 bucket
- Catalog
- Schema

In some cases, you might need to include a manifest file or a query.

- The user name that you want Amazon Quick Sight to use.

Every time Amazon Quick Sight connects using this data source profile (tile), it uses the user name from the connection settings. In some cases, this might be your personal login. But if you're going to share this with other people, ask the system administrator about creating credentials to use for Amazon Quick Sight connections.

- What type of connection to use. You can choose a public network or a VPC connection. If you have more than one VPC connection available, identify which one to use to reach your source of data.
- Additional settings, such as Secure Sockets Layer (SSL) or API tokens, are required by some data sources.

After you save the connection settings as a data source profile, you can create a dataset by selecting its tile. The connections are stored as data source connection profiles in Amazon Quick Sight.

To view your existing connection profiles, open the Quick start page, choose **Data**, choose **Create**, and then choose **New Dataset**.

For a list of supported data source connections and examples, see [Connect to your data with integrations and datasets](#).

After you create a data source in Quick Sight, you can [create a dataset](#) in Quick Sight that contains data from the connected data source. You can also [update data source connection](#) information at any time.

Editing a data source

You can edit an existing database data source to update the connection information, such as the server name or the user credentials. You can also edit an existing Amazon Athena data source to update the data source name. You can't edit Amazon S3 or Salesforce data sources.

Editing a database data source

Use the following procedure to edit a database data source.

1. From the Quick start page, choose **Data** at left. Choose **Create** and then choose **New dataset**.
2. Choose a database data source.
3. Choose **Edit Data Source**.
4. Modify the data source information:
 - If you are editing an autodiscovered database data source, you can modify any of the following settings:
 - For **Data source name**, enter a name for the data source.
 - For **Instance ID**, choose the name of the instance or cluster that you want to connect to from the list provided.
 - **Database name** shows the default database for the **Instance ID** cluster or instance. If you want to use a different database on that cluster or instance, enter its name.
 - For **UserName**, enter the user name of a user account that has permissions to do the following:
 - Access the target database.
 - Read (perform a SELECT statement on) any tables in that database that you want to use.
 - For **Password**, enter the password for the account that you entered.
 - If you are editing an external database data source, you can modify any of the following settings:
 - For **Data source name**, enter a name for the data source.
 - For **Database server**, enter one of the following values:

- For an Amazon Redshift cluster, enter the endpoint of the cluster without the port number. For example, if the endpoint value is `clustername.1234abcd.us-west-2.redshift.amazonaws.com:1234`, then enter `clustername.1234abcd.us-west-2.redshift.amazonaws.com`. You can get the endpoint value from the **Endpoint** field on the cluster detail page on the Amazon Redshift console.
 - For an Amazon EC2 instance of PostgreSQL, MySQL, or SQL Server, enter the public DNS address. You can get the public DNS value from the **Public DNS** field on the instance detail pane in the EC2 console.
 - For a non–Amazon EC2 instance of PostgreSQL, MySQL, or SQL Server, enter the hostname or public IP address of the database server.
 - For **Port**, enter the port that the cluster or instance uses for connections.
 - For **Database name**, enter the name of the database that you want to use.
 - For **UserName**, enter the user name of a user account that has permissions to do the following:
 - Access the target database.
 - Read (perform a SELECT statement on) any tables in that database that you want to use.
 - For **Password**, enter the password for the account that you entered.
5. Choose **Validate connection**.
 6. If the connection validates, choose **Update data source**. If not, correct the connection information and try validating again.
 7. If you want to create a new dataset using the updated data source, proceed with the instructions at [Creating a dataset from a database](#). Otherwise, close the **Choose your table** dialog box.

Editing an Athena data source

Use the following procedure to edit an Athena data source.

1. From the Quick start page, choose **Data** at left. Choose **Create** and then choose **New dataset**.
2. Choose an Athena data source.
3. Choose **Edit Data Source**.
4. For **Data source name**, enter a new name.

5. The **Manage data source sharing** screen appears. On the **Users** tab, locate the user that you want to remove.
6. If you want to create a new dataset using the updated data source, proceed with the instructions at [Creating a dataset using Amazon Athena data](#). Otherwise, close the **Choose your table** dialog box.

Deleting a data source

You can delete a data source if you no longer need it. Deleting a query-based database data source makes any associated datasets unusable. Deleting an Amazon S3, Salesforce, or SPICE-based database data source doesn't affect your ability to use any associated datasets. This is because the data is stored in [SPICE](#). However, you can no longer refresh those datasets.

To delete a data source

1. Choose the data source that you want to delete.
2. Choose **Delete**.

Refreshing data in Amazon Quick Sight

When refreshing data, Amazon Quick Sight handles datasets differently depending on the connection properties and the storage location of the data.

If Quick Sight connects to the data store by using a direct query, the data automatically refreshes when you open an associated dataset, analysis, or dashboard. Filter controls are refreshed automatically every 24 hours.

To refresh SPICE datasets, Quick Sight must independently authenticate using stored credentials to connect to the data. Quick Sight can't refresh manually uploaded data—even from S3 buckets, even though it's stored in SPICE—because Quick Sight doesn't store its connection and location metadata. If you want to automatically refresh data that's stored in an S3 bucket, create a dataset by using the **S3** data source card.

For files that you manually uploaded to SPICE, you refresh these manually by importing the file again. If you want to reuse the name of the original dataset for the new file, first rename or delete the original dataset. Then give the preferred name to the new dataset. Also, check that the field names are the same name and data type. Open your analysis, and replace the original dataset with the new dataset. For more information, see [Replacing datasets](#).

You can refresh your [SPICE](#) datasets at any time. Refreshing imports the data into SPICE again, so the data includes any changes since the last import.

For Amazon Quick Sight Standard Edition, you can do a full refresh of your SPICE data at any time. For Amazon Quick Sight Enterprise Edition, you can do a full refresh or an incremental refresh (SQL-based data sources only) at any time.

Note

If your dataset uses CustomSQL, refreshing incrementally might not benefit you. If the SQL query is complex, your database may not be able to optimize the filter with the look-back window. This can cause the query that pulls in the data to take longer than a full refresh. We recommend that you try reducing query execution time by refactoring the custom SQL. Note that results might vary depending on the type of optimization you make.

You can refresh SPICE data by using any of the following approaches:

- You can use the options on **Datasets** page.
- You can refresh a dataset while editing a dataset.
- You can schedule refreshes in the dataset settings.
- You can use the [CreateIngestion](#) API operation to refresh the data.

When you create or edit a SPICE dataset, you can enable email notifications about data loading status. This option notifies the owners of the dataset if the data fails to load or refresh. To turn on notifications, select the **Email owners when a refresh fails** option that appears on the **Finish dataset creation** screen. This option isn't available for datasets that you create by using **Upload a File** on the datasets page.

In the following topics, you can find an explanation of different approaches to refreshing and working with SPICE data.

Topics

- [Importing data into SPICE](#)
- [Refreshing SPICE data](#)
- [Using SPICE data in an analysis](#)
- [View SPICE ingestion history](#)

- [Troubleshooting skipped row errors](#)
- [SPICE ingestion error codes](#)
- [Updating files in a dataset](#)

Importing data into SPICE

When you import data into a dataset rather than using a direct SQL query, it becomes *SPICE data* because of how it's stored. *SPICE (Super-fast, Parallel, In-memory Calculation Engine)* is the robust in-memory engine that Amazon Quick Sight uses. It's engineered to rapidly perform advanced calculations and serve data. In Enterprise edition, data stored in SPICE is encrypted at rest.

When you create or edit a dataset, you choose to use either SPICE or a direct query, unless the dataset contains uploaded files. Importing (also called *ingesting*) your data into SPICE can save time and money:

- Your analytical queries process faster.
- You don't need to wait for a direct query to process.
- Data stored in SPICE can be reused multiple times without incurring additional costs. If you use a data source that charges per query, you're charged for querying the data when you first create the dataset and later when you refresh the dataset.

SPICE capacity is allocated separately for each AWS Region. Default SPICE capacity is automatically allocated to your home AWS Region. For each AWS account, SPICE capacity is shared by all the people using Quick Sight in a single AWS Region. The other AWS Regions have no SPICE capacity unless you choose to purchase some. Quick Sight administrators can view how much [SPICE](#) capacity you have in each AWS Region and how much of it is currently in use. A Quick Sight administrator can purchase more SPICE capacity or release unused SPICE capacity as needed. For more information, see [Configure SPICE memory capacity](#).

Topics

- [Estimating the size of SPICE datasets](#)

Estimating the size of SPICE datasets

The size of a dataset in SPICE relative to your Quick account's SPICE capacity is called *logical size*. A dataset's logical size isn't the same as the size of the dataset's source file or table. The computation

of a dataset's logical size occurs after all the data type transformations and calculated columns are defined during data preparation. These fields are materialized in SPICE in a way that enhances query performance. Any changes you make in an analysis have no effect on the logical size of the data in SPICE. Only changes that are saved in the dataset apply to SPICE capacity.

The logical size of a SPICE dataset depends on the data types of the dataset fields and the number of rows in the dataset. The three types of SPICE data are decimals, dates, and strings. You can transform a field's data type during the data preparation phase to fit your data visualization needs. For example, the file you want to import might contain all strings (text). But for these to be used in a meaningful way in an analysis, you prepare the data by changing the data types to their proper form. Fields containing prices can be changed from strings to decimals, and fields containing dates can be changed from strings to dates. You can also create calculated fields and exclude fields that you don't need from the source table. When you are finished preparing your dataset and all transformations are complete, you can estimate the logical size of the final schema.

 Note

Geospatial data types use metadata to interpret the physical data type. Latitude and longitude are numeric. All other geospatial categories are strings.

In the formula below, decimals and dates are calculated as 8 bytes per cell with 4 extra bytes for auxillary. Strings are calculated based on the text's length in UTF-8 encoding plus 24 bytes for auxillary. String data types require more space because of the extra indexing required by SPICE to provide high query performance.

```
Logical dataset size in bytes =  
(Number of Numeric cells * (12 bytes per cell))  
+ (Number of Date cells * (12 bytes per cell))  
+ SUM ((24 bytes + UTF-8 encoded length) per Text cell)
```

The formula above should only be used to estimate the size of a single dataset in SPICE. The SPICE capacity usage is the total size of all datasets in an account in a specific region. Quick Sight does not recommend that you use this formula to estimate the total SPICE capacity that your Quick Sight account is using.

Refreshing SPICE data

Refreshing a dataset

Use the following procedure to refresh a [SPICE](#) dataset based on an Amazon S3 or database data source in the **Data** tab. If there is a schema change in a database, Quick Sight will not be able to auto-detect it, resulting in an ingestion failure. Edit and save the dataset to update the schema and avoid ingestion failures.

To refresh SPICE data from the Data tab

1. Select **Data** from the left navigation menu. In the **Datasets** tab, choose the dataset to open it.
2. On the dataset details page that opens, choose the **Refresh** tab and then choose **Refresh now**.
3. Keep the refresh type as **Full refresh**.
4. If you are refreshing an Amazon S3 dataset, choose one of the following options for **S3 Manifest**:
 - To use the same manifest file you last provided to Amazon Quick Sight, choose **Existing Manifest**. If you have changed the manifest file at the file location or URL that you last provided, the data returned reflects those changes.
 - To specify a new manifest file by uploading it from your local network, choose **Upload Manifest**, and then choose **Upload manifest file**. For **Open**, choose a file, and then choose **Open**.
 - To specify a new manifest file by providing a URL, enter the URL of the manifest in **Input manifest URL**. You can find the manifest file URL in the Amazon S3 console by opening the context (right-click) menu for the manifest file, choosing **Properties**, and looking at the **Link** box.
5. Choose **Refresh**.
6. If you are refreshing an Amazon S3 dataset, choose **OK**, then **OK** again.

If you are refreshing a database dataset, choose **OK**.

Incrementally refreshing a dataset

Applies to: Enterprise Edition

For SQL-based data sources, such as Amazon Redshift, Amazon Athena, PostgreSQL, or Snowflake, you can refresh your data incrementally within a look-back window of time.

An *incremental refresh* queries only data defined by the dataset within a specified look-back window. It transfers all insertions, deletions, and modifications to the dataset, within that window's timeframe, from its source to the dataset. The data currently in SPICE that's within that window is deleted and replaced with the updates.

With incremental refreshes, less data is queried and transferred for each refresh. For example, let's say you have a dataset with 180,000 records that contains data from January 1 to June 30. On July 1, you run an incremental refresh on the data with a look-back window of seven days. Quick Sight queries the database asking for all data since June 24 (7 days ago), which is 7,000 records. Quick Sight then deletes the data currently in SPICE from June 24 and after, and appends the newly queried data. The next day (July 2), Quick Sight does the same thing, but queries from June 25 (7,000 records again), and then deletes from the existing dataset from the same date. Rather than having to ingest 180,000 records every day, it only has to ingest 7,000 records.

Use the following procedure to incrementally refresh a [SPICE](#) dataset based on a SQL data source from the **Datasets** tab.

To incrementally refresh a SQL-based SPICE dataset

1. Choose **Data** from the left navigation menu. On the **Datasets** tab, choose the dataset to open it.
2. On the dataset details page that opens, choose the **Refresh** tab and then choose **Refresh now**.
3. For **Refresh type**, choose **Incremental refresh**.
4. If this is your first incremental refresh on the dataset, choose **Configure**.
5. On the **Configure incremental refresh** page, do the following:
 - a. For **Date column**, choose a date column that you want to base the look-back window on.
 - b. For **Window size**, enter a number for **size**, and then choose an amount of time that you want to look back for changes.

You can choose to refresh changes to the data that occurred within a specified number of hours, days, or weeks from now. For example, you can choose to refresh changes to the data that occurred within two weeks of the current date.

6. Choose **Submit**.

Refreshing a dataset during data preparation

Use the following procedure to refresh a [SPICE](#) dataset based on an Amazon S3 or database data source during data preparation.

To refresh SPICE data during data preparation

1. Choose **Data** from the left navigation menu. On the **Datasets** tab, choose the dataset, and then choose **Edit Data Set**.
2. On the dataset screen, choose **Refresh now**.
3. Keep the refresh type set to **Full refresh**.
4. (Optional) If you are refreshing an Amazon S3 dataset, choose one of the following options for **S3 Manifest**:
 - To use the same manifest file that you last provided to Amazon Quick Sight, choose **Existing Manifest**. If you have changed the manifest file at the file location or URL that you last provided, the data returned reflects those changes.
 - To specify a new manifest file by uploading it from your local network, choose **Upload Manifest**, and then choose **Upload manifest file**. For **Open**, choose a file, and then choose **Open**.
 - To specify a new manifest file by providing a URL, enter the URL of the manifest in **Input manifest URL**. You can find the manifest file URL in the Amazon S3 console by opening the context (right-click) menu for the manifest file, choosing **Properties**, and looking at the **Link** box.
5. Choose **Refresh**.
6. If you are refreshing an Amazon S3 dataset, choose **OK**, then **OK** again.

If you are refreshing a database dataset, choose **OK**.

Refreshing a dataset on a schedule

Use the following procedure to schedule refreshing the data. If your dataset is based on a direct query and not stored in [SPICE](#), you can refresh your data by opening the dataset. You can also refresh your data by refreshing the page in an analysis or dashboard.

To refresh **SPICE** data on a schedule

1. Choose **Data** from the left navigation menu. On the **Datasets** tab, choose the dataset to open it.
 2. On the dataset details page that opens, choose the **Refresh** tab and then choose **Add new schedule**.
 3. On the **Create a refresh schedule** screen, choose settings for your schedule:
 - a. For **Time zone**, choose the time zone that applies to the data refresh.
 - b. For **Starting time**, choose a date and time for the refresh to start. Use HH:MM and 24-hour format, for example 13:30.
 - c. For **Frequency**, choose one of the following:
 - For Standard or Enterprise editions, you can choose **Daily**, **Weekly**, or **Monthly**.
 - **Daily**: Repeats every day.
 - **Weekly**: Repeats on the same day of each week.
 - **Monthly**: Repeats on the same day number of each month. To refresh data on the 29th, 30th or 31st day of the month, choose **Last day of month** from the list.
 - For Enterprise edition only, you can choose **Hourly**. This setting refreshes your dataset every hour, beginning at the time that you choose. So, if you select 1:05 as the starting time, the data refreshes at five minutes after the hour, every hour.
- If you decide to use an hourly refresh, you can't also use additional refresh schedules. To create an hourly schedule, remove any other existing schedules for that dataset. Also, remove any existing hourly schedule before you create a daily, weekly, or monthly schedule.
4. Choose **Save**.

Scheduled dataset ingestions take place within 10 minutes of the scheduled date and time.

Using the Quick console, you can create five schedules for each dataset. When you have created five, the **Create** button is turned off.

Incrementally refreshing a dataset on a schedule

Applies to: Enterprise Edition

For SQL-based data sources, such as Amazon Redshift, Athena, PostgreSQL, or Snowflake, you can schedule incremental refreshes. Use the following procedure to incrementally refresh a [SPICE](#) dataset based on a SQL data source in the **Datasets** tab.

To set an incremental refresh schedule for a SQL-based SPICE dataset

1. Choose **Data** from the left navigation menu. On the **Datasets** tab, choose the dataset to open it.
2. On the dataset details page that opens, choose the **Refresh** tab and then choose **Add new schedule**.
3. On the **Create a schedule** page, for **Refresh type**, choose **Incremental refresh**.
4. If this is your first incremental refresh for this dataset, choose **Configure**, and then do the following:
 - a. For **Date column**, choose a date column that you want to base the look-back window on.
 - b. For **Window size**, enter a number for **size**, and then choose an amount of time that you want to look back for changes.

You can choose to refresh changes to the data that occurred within a specified number of hours, days, or weeks from now. For example, you can choose to refresh changes to the data that occurred within two weeks of the current date.

- c. Choose **Submit**.
5. For **Time zone**, choose the time zone that applies to the data refresh.
 6. For **Repeats**, choose one of the following:
 - You can choose **Every 15 minutes**, **Every 30 minutes**, **Hourly**, **Daily**, **Weekly**, or **Monthly**.
 - **Every 15 minutes**: Repeats every 15 minutes, beginning at the time you choose. So, if you select 1:05 as the starting time, the data refreshes at 1:20, then again at 1:35, and so on.
 - **Every 30 minutes**: Repeats every 30 minutes, beginning at the time you choose. So, if you select 1:05 as the starting time, the data refreshes at 1:35, then again at 2:05, and so on.
 - **Hourly**: Repeats every hour, beginning at the time you choose. So, if you select 1:05 as the starting time, the data refreshes at five minutes after the hour, every hour.
 - **Daily**: Repeats every day.
 - **Weekly**: Repeats on the same day of each week.
 - **Monthly**: Repeats on the same day number of each month. To refresh data on the 29th, 30th or 31st day of the month, choose **Last day of month** from the list.

- If you decide to use refresh every 15 or 30 minutes, or hourly, you can't also use additional refresh schedules. To create a refresh schedule every 15 minutes, 30 minutes, or hourly, remove any other existing schedules for that dataset. Also, remove any existing minute or hourly schedule before you create a daily, weekly, or monthly schedule.
7. For **Starting**, choose a date for the refresh to start.
 8. For **At**, specify the time that the refresh should start. Use HH:MM and 24-hour format, for example 13:30.

Scheduled dataset ingestions take place within 10 minutes of the scheduled date and time.

In some cases, something might go wrong with the incremental refresh dataset that makes you want to roll back your dataset. Or you might no longer want to refresh the dataset incrementally. If so, you can delete the scheduled refresh.

To do so, choose the dataset on the **Datasets** page, choose **Schedule a refresh**, and then choose the x icon to the right of the scheduled refresh. Deleting an incremental refresh configuration starts a full refresh. As part of this full refresh, all the configurations prepared for incremental refreshes are removed.

Using SPICE data in an analysis

When you use stored data to create an analysis, a data import indicator appears next to the dataset list at the top of the **Fields list** pane. When you first open the analysis and the dataset is importing, a spinner icon appears.

After the SPICE import is complete, the indicator displays the percentage of rows that were successfully imported. A message also appears at the top of the visualization pane to provide counts of the rows imported and skipped.

If any rows were skipped, you can choose **View summary** in this message bar to see details about why those rows failed to import. To edit the dataset and resolve the issues that led to skipped rows, choose **Edit data set**. For more information about common causes for skipped rows, see [Troubleshooting skipped row errors](#).

If an import fails altogether, the data import indicator appears as an exclamation point icon, and an **Import failed** message is displayed.

View SPICE ingestion history

You can view the ingestion history for SPICE datasets to find out, for example, when the latest ingestion started and what its status is.

The SPICE ingestion history page includes the following information:

- Date and time that the ingestion started (UTC)
- Status of the ingestion
- Amount of time that the ingestion took
- The number of aggregated rows in the dataset.
- The number of rows ingested during a refresh.
- Rows skipped and rows ingested (imported) successfully
- The job type for the refresh: scheduled, full refresh, and so on

Use the following procedure to view a dataset's SPICE ingestion history.

To view a dataset's SPICE ingestion history

1. From the homepage, choose **Data** at left.
2. On the **Datasets** tab, choose the dataset that you want to examine.
3. On the dataset details page that opens, choose the **Refresh** tab.

SPICE ingestion history is shown at bottom.

4. (Optional) Choose a time frame to filter the entries from the last hour to the last 90 days.
5. (Optional) Choose a specific job status to filter the entries, for example **Running** or **Completed**. Otherwise, you can view all entries by choosing **All**.

Troubleshooting skipped row errors

When you import data, Amazon Quick Sight previews a portion of your data. If it can't interpret a row for any reason, Quick Sight skips the row. In some cases, the import will fail. When this happens, Quick Sight returns an error message that explains the failure.

Fortunately, there's a limited number of things that can go wrong. Some issues can be avoided by being aware of examples like the following:

- Make sure that there is no inconsistency between the field data type and the field data, for example occasional string data in a field with a numeric data type. Here are a few examples that can be difficult to detect when scanning the contents of a table:
 - ' ' – Using an empty string to indicate a missing value
 - 'NULL ' – Using the word "null" to indicate a missing value
 - \$1000 – Including a dollar sign in a currency value turns it into a string
 - 'O'Brien' – Using punctuation to mark a string that itself contains the same punctuation.

However, this type of error isn't always this easy to find, especially if you have a lot of data, or if your data is typed in by hand. For example, some customer service or sales applications involve entering information verbally provided by customers. The person who originally typed in the data might have put it in the wrong field. They might add, or forget to add, a character or digit. For example, they might enter a date of "0/10/12020" or enter someone's gender in a field meant for age.

- Make sure that your imported file is correctly processed with or without a header. If there is a header row, make sure that you choose the **Contains header** upload option.
- Make sure that the data doesn't exceed one or more of the [Data source quotas](#).
- Make sure that the data is compatible with the [Supported data types and values](#).
- Make sure that your calculated fields contain data that works with the calculation, rather than being incompatible with or excluded by the function in the calculated field. For example, if you have a calculated field in your dataset that uses [parseDate](#), Quick Sight skips rows where that field doesn't contain a valid date.

Quick Sight provides a detailed list of the errors that occur when the SPICE engine attempts to ingest data. When a saved dataset reports skipped rows, you can view the errors so you can take action to fix the issues.

To view errors for rows that were skipped during SPICE ingestion (data import)

1. Choose **Data** at left. In the **Datasets** tab, choose the problematic dataset to open it.
2. On the dataset details page that opens, choose the **Refresh** tab.

SPICE ingestion history is shown at bottom.

3. For the ingestion with the error, choose **View error summary**. This link is located under the **Status** column.

4. Examine the **File import log** that opens. It displays the following sections:
 - **Summary** – Provides a percentage score of how many rows were skipped out of the total number of rows in the import. For example, if there are 864 rows skipped out of a total of 1,728, the score is 50.00%.
 - **Skipped Rows** – Provides the row count, field name, and error message for each set of similar skipped rows.
 - **Troubleshooting** – Provides a link to download a file that contains error information.
5. Under **Troubleshooting**, choose **Download error rows file**.

The error file has a row for each error. The file is named `error-report_123_fe8.csv`, where `123_fe8` is replaced with a unique identifying string. The file contains the following columns:

- **ERROR_TYPE** – The type or error code for the error that occurred when importing this row. You can look up this error in the [SPICE ingestion error codes](#) section that follows this procedure.
 - **COLUMN_NAME** – The name of the column in your data that caused the error.
 - All the columns from your imported row – The remaining columns duplicate the entire row of data. If a row has more than one error, it can appear multiple times in this file.
6. Choose **Edit data set** to make changes to your dataset. You can filter the data, omit fields, change data types, adjust existing calculated fields, and add calculated fields that validate the data.
 7. After you've made changes indicated by the error codes, import the data again. If more SPICE ingestion errors appear in the log, step through this procedure again to fix all remaining errors.

Tip

If you can't solve the data issues in a reasonable amount of time by using the dataset editor, consult the administrators or developers who own the data. In the long run, it's more cost-effective to cleanse the data closer to its source, rather than adding exception processing while you're preparing the data for analysis. By fixing it at the source, you avoid a situation where multiple people fix the errors in different ways, resulting in different reporting results later on.

To practice troubleshooting skipped rows

1. Download [CSV files for troubleshooting skipped rows.zip](#).
2. Extract the files into a folder that you can use to upload the sample .csv file into Quick Sight.

The zip file contains the following two text files:

- `sample dataset - data ingestion error.csv` – A sample .csv file that contains issues that cause skipped rows. You can try to import the file yourself to see how the error process works.
 - `sample data ingestion error file` – A sample error file generated during SPICE ingestion while importing the sample .csv file into Quick Sight.
3. Import the data by following these steps:
 - a. Choose **Data, Datasets** tab, **New, Dataset**.
 - b. Choose **Upload a file**.
 - c. Find and choose the file named `sample dataset - data ingestion error.csv`.
 - d. Choose **Upload a file, Edit settings and prepare data**.
 - e. Choose **Save** to exit.
 4. Choose your dataset to view its information, then choose **View error summary**. Examine the errors and the data to help you resolve the issues.

SPICE ingestion error codes

The following list of errors codes and descriptions can help you understand and troubleshoot issues with data ingestion into SPICE.

Error codes for skipped rows

The following list of errors codes and descriptions can help you understand and troubleshoot issues with skipped rows.

ARITHMETIC_EXCEPTION – An arithmetic exception occurred while processing a value.

ENCODING_EXCEPTION – An unknown exception occurred while converting and encoding data to SPICE.

OPENSEARCH_CURSOR_NOT_ENABLED – The OpenSearch domain doesn't have SQL cursors enabled ("opendistro.sql.cursor.enabled" : "true"). For more information, see [Authorizing connections to Amazon OpenSearch Service](#).

INCORRECT_FIELD_COUNT – One or more rows have too many fields. Make sure that the number of fields in each row matches the number of fields defined in the schema.

INCORRECT_SAGEMAKER_OUTPUT_FIELD_COUNT – The SageMaker AI output has an unexpected number of fields.

INDEX_OUT_OF_BOUNDS – The system requested an index that isn't valid for the array or list being processed.

MALFORMED_DATE – A value in a field can't be transformed to a valid date. For example, if you try to convert a field that contains a value like "sale date" or "month-1", the action generates a malformed date error. To fix this error, remove nondate values from your data source. Check that you aren't importing a file with a column header mixed into the data. If your string contains a date or time that doesn't convert, see [Using unsupported or custom dates](#).

MISSING_SAGEMAKER_OUTPUT_FIELD – A field in the SageMaker AI output is unexpectedly empty.

NUMBER_BITWIDTH_TOO_LARGE – A numeric value exceeds the length supported in SPICE. For example, your numeric value has more than 19 digits, which is the length of a bigint data type. For a long numeric sequence that isn't a mathematical value, use a string data type.

NUMBER_PARSE_FAILURE – A value in a numeric field is not a number. For example, a field with a data type of int contains a string or a float.

SAGEMAKER_OUTPUT_COLUMN_TYPE_MISMATCH – The data type defined in the SageMaker AI schema doesn't match the data type received from SageMaker AI.

STRING_TRUNCATION – A string is being truncated by SPICE. Strings are truncated where the length of the string exceeds the SPICE quota. For more information about SPICE, see [Importing data into SPICE](#). For more information about quotas, see [Service Quotas](#).

UNDEFINED – An unknown error occurred while ingesting data.

UNSUPPORTED_DATE_VALUE – A date field contains a date that is in a supported format but is not in the supported range of dates, for example "12/31/1399" or "01/01/10000". For more information, see [Using unsupported or custom dates](#).

Error codes during data import

For imports and data refresh jobs that fail, Quick Sight provides an error code indicating what caused the failure. The following list of errors codes and descriptions can help you understand and troubleshoot issues with data ingestion into SPICE.

ACCOUNT_CAPACITY_LIMIT_EXCEEDED – This data exceeds your current SPICE capacity. Purchase more SPICE capacity or clean up existing SPICE data and then retry this ingestion.

CONNECTION_FAILURE – Amazon Quick Sight can't connect to your data source. Check the data source connection settings and try again.

CUSTOMER_ERROR – There was a problem parsing the data. If this persists, contact Amazon Quick Sight technical support.

DATA_SET_DELETED – The data source or dataset was deleted or became unavailable during ingestion.

DATA_SET_SIZE_LIMIT_EXCEEDED – This dataset exceeds the maximum allowable SPICE dataset size. Use filters to reduce the dataset size and try again. For information on SPICE quotas, see [Data source quotas](#).

DATA_SOURCE_AUTH_FAILED – Data source authentication failed. Check your credentials and use the **Edit data source** option to replace expired credentials.

DATA_SOURCE_CONNECTION_FAILED – Data source connection failed. Check the URL and try again. If this error persists, contact your data source administrator for assistance.

DATA_SOURCE_NOT_FOUND – No data source found. Check your Amazon Quick Sight data sources.

DATA_TOLERANCE_EXCEPTION – There are too many invalid rows. Amazon Quick Sight has reached the quota of rows that it can skip and still continue ingesting. Check your data and try again.

FAILURE_TO_ASSUME_ROLE – Amazon Quick Sight couldn't assume the correct AWS Identity and Access Management (IAM) role. Verify the policies for Amazon Quick Sight-service-role in the IAM console.

FAILURE_TO_PROCESS_JSON_FILE – Amazon Quick Sight couldn't parse a manifest file as valid JSON.

IAM_ROLE_NOT_AVAILABLE – Amazon Quick Sight doesn't have permission to access the data source. To manage Amazon Quick Sight permissions on AWS resources, go to the **Security and Permissions** page under the **Manage Amazon Quick Sight** option as an administrator.

INGESTION_CANCELED – The ingestion was canceled by the user.

INGESTION_SUPERSEDED – This ingestion has been superseded by another workflow. This happens when a new ingestion is created while another one is still in progress. Avoid manually editing the dataset multiple times in a short period, because each manual edit creates a new ingestion which will supersede and end the previous ingestion.

INTERNAL_SERVICE_ERROR – An internal service error occurred.

INVALID_DATA_SOURCE_CONFIG – Invalid values appeared in connection settings. Check your connection details and try again.

INVALID_DATAPREP_SYNTAX – Your calculated field expression contains invalid syntax. Correct the syntax and try again.

INVALID_DATE_FORMAT – An invalid date format appeared.

IOT_DATA_SET_FILE_EMPTY – No AWS IoT Analytics data was found. Check your account and try again.

IOT_FILE_NOT_FOUND – An indicated AWS IoT Analytics file wasn't found. Check your account and try again.

OAuth_TOKEN_FAILURE – Credentials to the data source have expired. Renew your credentials and retry this ingestion.

PASSWORD_AUTHENTICATION_FAILURE – Incorrect credentials appeared for a data source. Update your data source credentials and retry this ingestion.

PERMISSION_DENIED – Access to the requested resources was denied by the data source. Request permissions from your database administrator or ensure proper permission has been granted to Amazon Quick Sight before retrying.

QUERY_TIMEOUT – A query to the data source timed out waiting for a response. Check your data source logs and try again.

ROW_SIZE_LIMIT_EXCEEDED – The row size quota exceeded the maximum.

S3_FILE_INACCESSIBLE – Couldn't connect to an S3 bucket. Make sure that you grant Amazon Quick Sight and users necessary permissions before you connect to the S3 bucket.

S3_MANIFEST_ERROR – Couldn't connect to S3 data. Make sure that your S3 manifest file is valid. Also verify access to the S3 data. Both Amazon Quick Sight and the Amazon Quick Sight user need permissions to connect to the S3 data.

S3_UPLOADED_FILE_DELETED – The file or files for the ingestion were deleted (between ingestions). Check your S3 bucket and try again.

SOURCE_API_LIMIT_EXCEEDED_FAILURE – This ingestion exceeds the API quota for this data source. Contact your data source administrator for assistance.

SOURCE_RESOURCE_LIMIT_EXCEEDED – A SQL query exceeds the resource quota of the data source. Examples of resources involved can include the concurrent query quota, the connection quota, and physical server resources. Contact your data source administrator for assistance.

SPICE_TABLE_NOT_FOUND – An Amazon Quick Sight data source or dataset was deleted or became unavailable during ingestion. Check your dataset in Amazon Quick Sight and try again. For more information, see [Troubleshooting skipped row errors](#).

SQL_EXCEPTION – A general SQL error occurred. This error can be caused by query timeouts, resource constraints, unexpected data definition language (DDL) changes before or during a query, and other database errors. Check your database settings and your query, and try again.

SQL_INVALID_PARAMETER_VALUE – An invalid SQL parameter appeared. Check your SQL and try again.

SQL_NUMERIC_OVERFLOW – Amazon Quick Sight encountered an out-of-range numeric exception. Check related values and calculated columns for overflows, and try again.

SQL_SCHEMA_MISMATCH_ERROR – The data source schema doesn't match the Amazon Quick Sight dataset. Update your Amazon Quick Sight dataset definition.

SQL_TABLE_NOT_FOUND – Amazon Quick Sight can't find the table in the data source. Verify the table specified in the dataset or custom SQL and try again.

SSL_CERTIFICATE_VALIDATION_FAILURE – Amazon Quick Sight can't validate the Secure Sockets Layer (SSL) certificate on your database server. Check the SSL status on that server with your database administrator and try again.

UNRESOLVABLE_HOST – Amazon Quick Sight can't resolve the host name of the data source. Verify the host name of the data source and try again.

UNROUTABLE_HOST – Amazon Quick Sight can't reach your data source because it's inside a private network. Ensure that your private VPC connection is configured correctly in Enterprise Edition, or allow Amazon Quick Sight IP address ranges to allow connectivity for Standard Edition.

Updating files in a dataset

To get the latest version of files, you can update files in your dataset. You can update these types of files:

- Comma-delimited (CSV) and tab-delimited (TSV) text files
- Extended and common log format files (ELF and CLF)
- Flat or semistructured data files (JSON)
- Microsoft Excel (XLSX) files

Before updating a file, make sure that the new file has the same fields in the same order as the original file currently in the dataset. If there are field (column) discrepancies between the two files, an error occurs and you need to fix the discrepancies before attempting to update again. You can do this by editing the new file to match the original. Note that if you want to add new fields, you can append them after the original fields in the file. For example, in a Microsoft Excel spreadsheet, you can append new fields to the right of the original fields.

To update a file in a dataset

1. In Quick Sight, choose **Data** at left.
2. In the **Datasets** tab, choose the dataset that you want to update, and then choose **Edit dataset**.
3. On the data preparation page that opens, choose the drop-down list for the file that you want to update, and then choose **Update file**.
4. On the **Update file** page that opens, choose **Upload file**, and then navigate to a file.

Quick Sight scans the file.

5. If the file is a Microsoft Excel file, choose the sheet that you want on the **Choose your sheet** page that opens, and then choose **Select**.

6. Choose **Confirm file update** on the following page. A preview of some of the sheet columns is shown for your reference.

A message that the file updated successfully appears at top right and the table preview updates to show the new file data.

Preparing data in Amazon Quick Sight

Datasets store any data preparation you have done on that data, so that you can reuse that prepared data in multiple analyses. Data preparation provides options such as adding calculated fields, applying filters, and changing field names or data types. If you are basing the data source on a SQL database, you can also use data preparation to join tables. Or you can enter a SQL query if you want to work with data from more than a single table.

If you want to transform the data from a data source before using it in Amazon Quick Sight, you can prepare it to suit your needs. You then save this preparation as part of the dataset.

You can prepare a dataset when you create it, or by editing it later. For more information about creating a new dataset and preparing it, see [Creating datasets](#). For more information about opening an existing dataset for data preparation, see [Editing datasets](#).

Use the following topics to learn more about data preparation.

Topics

- [Data Preparation Experience \(New\)](#)
- [Describing data](#)
- [Choosing file upload settings](#)
- [Data Preparation Experience \(Legacy\)](#)
- [Using SQL to customize data](#)
- [Adding geospatial data](#)
- [Using unsupported or custom dates](#)
- [Adding a unique key to an Amazon Quick Sight dataset](#)
- [Integrating Amazon SageMaker AI models with Amazon Quick Sight](#)
- [Preparing dataset examples](#)

Data Preparation Experience (New)

Data preparation transforms raw data into a format optimized for analysis and visualization. In business intelligence, this crucial process involves cleaning, structuring, and enriching data to enable meaningful business insights.

Amazon Quick Sight's data preparation interface revolutionizes this process with an intuitive, visual experience that enables users to create analysis-ready datasets without SQL expertise. Through its modern, streamlined approach, users can efficiently create and manage business intelligence datasets. The visual interface presents a clear, sequential view of data transformations, allowing authors to track changes from the initial state to the final output with precision.

The platform emphasizes collaboration and reusability, enabling teams to share and repurpose workflows across the organization. This collaborative design promotes consistency in data transformation practices while eliminating redundant work, ultimately fostering standardized processes across teams and enhancing overall efficiency.

Topics

- [Components within the data preparation experience](#)
- [Data preparation steps](#)
- [Advanced workflow capabilities](#)
- [SPICE-only features](#)
- [Switching between data preparation experiences](#)
- [Features not supported in the new data preparation experience](#)
- [Data preparation limits](#)
- [Ingestion behavior changes](#)
- [Frequently asked questions](#)
- [Dataset Enrichment](#)

Components within the data preparation experience

Amazon Quick Sight's data preparation experience has the following core components.

Workflow

A workflow in Quick Sight's data preparation experience represents a sequential series of data transformation steps that guide your dataset from its raw state to an analysis-ready form. These

workflows are designed for reusability, enabling analysts to leverage and build upon existing work while maintaining consistent data transformation standards throughout the organization.

While workflows can accommodate multiple paths through various Inputs or through Divergence (detailed in subsequent sections), they must ultimately converge into a single output table. This unified structure ensures data consistency and streamlined analysis capabilities.

Transformation

A transformation is a specific data manipulation operation that changes the structure, format, or content of your data. Quick Sight's data preparation experience offers various transformation types including Join, Filter, Aggregate, Pivot, Unpivot, Append, and Calculated Columns.

Each transformation type serves a distinct purpose in reshaping your data to meet analytical requirements. These transformations are implemented as individual steps within your workflow.

Step

A step is a collection of homogeneous transformations of the same type applied within your workflow. Each step contains one or more related operations of the same transformation category. For example, a Rename step can include multiple column renaming operations, and a Filter step can contain multiple filtering conditions—all managed as a single unit in your workflow.

Most steps can include multiple operations, with two notable exceptions: Join and Append steps are limited to two input tables per step. To join or append more than two tables, you can create additional Join or Append steps in sequence.

Steps are displayed in order, with each step building upon the results of previous steps, allowing you to track the progressive transformation of your data. To rename or delete a step, select it and choose the three-dot menu.

Connector

The connector links two steps with an arrow indicating the workflow direction. You can delete a connector by selecting it and pressing the delete key. To add a step between two existing steps, simply delete the connector, add the new step, and reconnect the steps by dragging your mouse between them.

Configure pane

The **Configuration pane** is the interactive area where you define parameters and settings for a selected step. When you select a step in your workflow, this pane displays relevant options for

that specific transformation type. For example, when configuring a Join step, you can select join type, matching columns, and other join-specific settings. The **Configuration pane**'s point-and-click interface eliminates the need for SQL knowledge.

Preview pane

The **Preview pane** displays a real-time sample of your data as it appears after applying the current transformation step. This immediate visual feedback helps you verify that each transformation produces the expected results before proceeding to the next step. The **Preview pane** updates dynamically as you modify step configurations, enabling iterative refinement of data transformations with confidence.

These components work together to create an intuitive, visual data preparation experience that makes complex data transformations accessible to business users without requiring technical expertise.

Data preparation steps

Amazon Quick Sight's data preparation experience offers eleven powerful step types that enable you to transform your data systematically. Each step serves a specific purpose in the data preparation workflow.

Steps can be configured through an intuitive interface in the **Configuration pane**, with immediate feedback visible in the **Preview pane**. Steps can be combined sequentially to create sophisticated data transformations without requiring SQL expertise.

Each step can receive input from either a physical table or the output of a previous step. Most steps accept a single input, with Append and Join steps being the exceptions—these require exactly two inputs.

Input

The Input step initiates your data preparation workflow in Quick Sight by allowing you to select and import data from multiple sources for transformation in subsequent steps.

Input options

- **Add Dataset**

Leverage existing Quick Sight datasets as input sources, building upon data that has already been prepared and optimized by your team.

- **Add Data Source**

Connect directly to databases such as Amazon Redshift, Athena, RDS, or other supported sources by selecting specific database objects and providing connection parameters.

- **Add File Upload**

Import data directly from local files in formats such as CSV, TSV, Excel, or JSON.

Configuration

The Input step requires no configuration. The **Preview** pane displays your imported data along with source information, including connection details, table name, and column metadata.

Usage notes

- Multiple Input steps can exist within a single workflow.
- You can add Input steps at any point in your workflow.

Add Calculated Columns

The Add Calculated Columns step enables you to create new columns using row-level expressions that perform calculations on existing columns. You can create new columns using scalar (row-level) functions and operators, and apply row-level calculations that reference existing columns.

Configuration

To configure the Add Calculated Columns step, in the **Configuration** pane:

1. Name your new calculated column.
2. Build expressions using the calculation editor, which supports row-level functions and operators (such as [ifelse](#) and [round](#)).
3. Save your calculation.
4. Preview the expression results.
5. Add more calculated columns as needed.

Usage notes

- Only scalar (row-level) calculations are supported in this step.

- In SPICE, calculated columns are materialized and function as standard columns in subsequent steps.

Change Data Type

Quick Sight simplifies data type management by supporting four abstract data types: `date`, `decimal`, `integer`, and `string`. These abstract types eliminate complexity by automatically mapping various source data types to their Quick Sight equivalents. For instance, `tinyint`, `smallint`, `integer`, and `bigint` are all mapped to `integer`, while `date`, `datetime`, and `timestamp` are mapped to `date`.

This abstraction means you only need to understand Quick Sight's four data types, as Quick Sight handles all underlying data type conversions and calculations automatically when interacting with different data sources.

Configuration

To configure the Change Data Type step, in the **Configuration** pane:

1. Select a column to convert.
2. Choose the target data type (`string`, `integer`, `decimal`, or `date`).
3. For date conversions, specify format settings and preview results based on input formats. See the [supported date formats](#) in Quick Sight.
4. Add additional columns to convert as needed.

Usage notes

- Convert multiple columns' data types in a single step for efficiency.
- When using SPICE, all data type changes are materialized in the imported data.

Rename Columns

The Rename Columns step enables you to modify column names to be more descriptive, user-friendly, and consistent with your organization's naming conventions.

Configuration

To configure the Rename Columns step, in the **Configuration** pane:

1. Select a column to name.
2. Enter a new name for the selected column.
3. Add more columns to rename as needed.

Usage notes

- All column names must be unique within your dataset.

Select Columns

The Select Columns step enables you to streamline your dataset by including, excluding, and reordering columns. This helps optimize your data structure by removing unnecessary columns and organizing the remaining ones in a logical sequence for analysis.

Configuration

To configure the Select Columns step, in the **Configuration** pane:

1. Choose specific columns to include in your output.
2. Select columns in your preferred order to establish sequence.
3. Use **Select All** to include remaining columns in their original order.
4. Exclude unwanted columns by leaving them unselected.

Key Features

- Output columns appear in the order of selection.
- **Select All** preserves the original column sequence.

Usage notes

- Unselected columns are removed from subsequent steps.
- Optimize dataset size by removing unnecessary columns.

Append

The Append step vertically combines two tables, similar to a SQL UNION ALL operation. Quick Sight automatically matches columns by name rather than sequence, enabling efficient data consolidation even when tables have different column orders or varying numbers of columns.

Configuration

To configure the Append step, in the **Configuration** pane:

1. Select two input tables to append.
2. Review the output column sequence.
3. Examine which columns are present in both tables versus single tables.

Key features

- Matches columns by name instead of sequence.
- Retains all rows from both tables, including duplicates.
- Supports tables with different numbers of columns.
- Follows Table 1's column sequence for matching columns, then adds unique columns from Table 2.
- Shows clear source indicators for all columns

Usage notes

- Use a Rename step first when appending columns with different names.
- Each Append step combines exactly two tables; use additional Append steps for more tables.

Join

The Join step horizontally combines data from two tables based on matching values in specified columns. Quick Sight supports Left Outer, Right Outer, Full Outer, and Inner Join types, providing flexible options for your analytical needs. The step includes intelligent column conflict resolution that automatically handles duplicate column names. While self-joins aren't available as a specific join type, you can achieve similar results using workflow divergence.

Configuration

To configure the Join step, in the **Configuration** pane:

1. Select two input tables to join.
2. Choose your join type (Left Outer, Right Outer, Full Outer, or Inner).
3. Specify join keys from each table.
4. Review auto-resolved column name conflicts.

Key features

- Supports multiple join types for different analytical needs.
- Automatically resolves duplicate column names.
- Accepts calculated columns as join keys.

Usage notes

- Join keys must have compatible data types; use the Change Data Type step if needed.
- Each Join step combines exactly two tables; use additional Join steps for more tables.
- Create a Rename step after the Join to customize auto-resolved column headers.

Aggregate

The Aggregate step enables you to summarize data by grouping columns and applying aggregation operations. This powerful transformation condenses detailed data into meaningful summaries based on your specified dimensions. Quick Sight simplifies complex SQL operations through an intuitive interface, offering comprehensive aggregation functions including advanced string operations like `ListAgg` and `ListAgg distinct`.

Configuration

To configure the Aggregate step, in the **Configuration** pane:

1. Select columns to group by.
2. Choose aggregation functions for measure columns.
3. Customize output column names.
4. For `ListAgg` and `ListAgg distinct`:

- a. Select the column to aggregate.
 - b. Choose a separator (comma, dash, semicolon, or vertical line).
5. Preview the summarized data.

Supported functions per data type

Data Type	Supported Functions
Numeric	Average, Sum Count, Count Distinct Max, Min
Date	Count, Count Distinct Max, Min ListAgg, ListAgg distinct (for date only)
String	ListAgg, ListAgg distinct Count, Count Distinct Max, Min

Key features

- Applies different aggregation functions to columns within the same step.
- **Group by** without aggregation functions acts as SQL SELECT DISTINCT.
- ListAgg concatenates all values; ListAgg distinct includes only unique values.
- ListAgg functions maintain ascending sort order by default.

Usage notes

- Aggregation significantly reduces row count in your dataset.
- ListAgg and ListAgg distinct support date values but not datetime.
- Use separators to customize string concatenation output.

Filter

The Filter step enables you to narrow down your dataset by including only rows that meet specific criteria. You can apply multiple filter conditions within a single step, all combining through AND logic to help focus your analysis on relevant data.

Configuration

To configure the Filter step, in the **Configuration** pane:

1. Select a column to filter.
2. Choose a comparison operator.
3. Specify filter values based on the column's data type.
4. Add additional filter conditions across different columns if needed.

Note

- String filters with "is in" or "is not in": Enter multiple values (one per line).
- Numeric and date filters: Enter single values (except "between" which requires two values).

Supported operators per data type

Data Type	Supported Operators
Integer and Decimal	Equals, Does not equal
	Greater than, Less than
	Is greater than or equal to, Is less than or equal to

Data Type	Supported Operators
	Is between
Date	After, Before Is between Is after or equal to, Is before or equal to
String	Equals, Does not equal Starts with, Ends with Contains, Does not contain Is in, Is not in

Usage notes

- Apply multiple filter conditions in a single step.
- Mix conditions across different data types.
- Preview filtered results in real-time.

Pivot

The Pivot step transforms row values into unique columns, converting data from a long format to a wide format for easier comparison and analysis. This transformation requires specifications for value filtering, aggregation, and grouping to manage the output columns effectively.

Configuration

To configure the Pivot step, use the following in the **Configuration** pane:

1. **Pivot column:** Select the column whose values will become column headers (e.g., Category).
2. **Pivot column row value:** Filter specific values to include (e.g., Technology, Office Supplies).
3. **Output column header:** Customize new column headers (defaults to pivot column values).

4. **Value column:** Select the column to aggregate (e.g., Sales).
5. **Aggregation function:** Choose the aggregation method (e.g., Sum).
6. **Group by:** Specify organizing columns (e.g., Segment).

Input table (full table not shown):

product_name	ship_mode	sales
Portfolio Organizer	Same Day	3054.32
USB Hub	Next Day	3465.1
Mesh Task Chair	Express	3051.43
Wall Clock	Express	3974.86
Storage Box	Next Day	321.69
Archive System	Express	644.26
Archive System	Same Day	32.14
Work Table	Same Day	1030.33
Wall Clock	Standard	568.31

Output table (full table not shown):

product_name	Standard Sales	Same Day Sales	Next Day Sales
Wireless Headset	41244.5	13722.52	21537.18
Professional Art Set	30901.940000000002	17425.24	20884.13
Ergonomic Office Chair	24058.309999999998	26860.940000000002	51962.05
View Binder	30875.64	17954.690000000002	15480.54
Drawing Markers	41170.240000000005	14449.169999999998	15397.899999999999
Wall Clock	11756.699999999999	20273.260000000002	32238.179999999997
Whiteboard	12253.109999999999	18574.629999999997	31560.699999999997
Label Printer	14349.65	18759.739999999998	12415.79

Configuration:

Pivot column

Column values become new column headers in output

Pivot column

ship_mode

Only STRING data type is supported

Pivot column row value

Standard

Same Day

Next Day

Output column header

Standard Sales

Same Day Sales

Next Day Sales

Actions

Value column

Column values fill rows of pivot output columns

Value column

sales

Aggregation function

Sum

Group by columns (optional)

Columns identify rows of output

1 group by column

Group by columns

product_name

Supported operators per data type

Data Type	Supported Operators
Integer and Decimal	Average, Sum Count, Count Distinct Max, Min
Date	Count, Count Distinct Max, Min ListAgg, ListAgg distinct (date values only)
String	

Data Type	Supported Operators
	ListAgg, ListAgg distinct
	Count, Count Distinct
	Max, Min

Usage notes

- Each pivoted column contains aggregated values from the value column.
- Customize column headers for clarity.
- Preview transformation results in real-time.

Unpivot

The Unpivot step transforms columns into rows, converting wide data into a longer, narrower format. This transformation helps organize data spread across multiple columns into a more structured format for easier analysis and visualization.

Configuration

To configure the Unpivot step, in the **Configuration** pane:

1. Select columns to unpivot into rows.
2. Define output column row values. The default is the original column name. Some examples include Technology, Office Supplies, and Furniture.
3. Name the two new outputs columns.
 - **Unpivoted column header:** The name for former column names (e.g., Category)
 - **Unpivoted column values:** The name for the unpivoted values (e.g., Sales)

Input table (full table not shown):

product_name	Standard Sales	Same Day Sales	Next Day Sales
Wireless Headset	42444.6	15522.62	21557.18
Professional Art Set	50901.940000000002	17425.24	20984.13
Ergonomic Office Chair	24058.309999999998	26860.940000000002	51962.05
View Binder	30875.54	17954.690000000002	15480.54
Drawing Markers	41170.240000000005	16449.169999999998	15387.789999999999
Wall Clock	15756.699999999999	20275.260000000002	32238.179999999997
Whiteboard	12253.109999999999	18574.629999999997	51560.699999999997
Label Printer	14349.65	18759.739999999998	12415.79
Laser Printer	46872.649999999994	29048.97	10125.82
Executive Leather Chair	23222.37	11613.26	14295.25

Output table (full table not shown):

product_name	Ship Mode	Sales
Wireless Headset	Standard	42444.6
Professional Art Set	Standard	50901.940000000002
Ergonomic Office Chair	Standard	24058.309999999998
View Binder	Standard	30875.54
Drawing Markers	Standard	41170.240000000005
Wall Clock	Standard	15756.699999999999
Whiteboard	Standard	12253.109999999999
Label Printer	Standard	14349.65
Laser Printer	Standard	46872.649999999994

Configuration:

Unpivot columns

Select columns to unpivot

product_name

Standard Sales

Same Day Sales

Next Day Sales

Unpivot 1 columns

Input column

Standard Sales

Same Day Sales

Next Day Sales

Output column row value

Standard

Same Day

Next Day

Actions

Name output columns

Columns for unpivoted column headers

Ship Mode

Columns for unpivoted column values

Sales

Key features

- Retains all non-unpivoted columns in the output.
- Creates two new columns automatically: one for former column names and one for their corresponding values.
- Transforms wide data into long format.

Usage notes

- All unpivoted columns must have compatible data types.
- Row count typically increases after unpivoting.
- Preview changes in real-time before applying them.

Advanced workflow capabilities

Amazon Quick Sight's data preparation experience offers sophisticated features that enhance your ability to create complex, reusable data transformations. This section covers two powerful capabilities that extend your workflow potential.

Divergence enables you to create multiple transformation paths from a single step, allowing parallel processing streams that can be recombined later. This capability is particularly valuable for complex scenarios like self-joins and parallel transformations.

Composite Datasets allow you to build hierarchical data structures by using existing datasets as building blocks. This feature promotes collaboration across teams and ensures consistent business logic through reusable, layered transformations.

These capabilities work together to provide flexible workflow designs, enhanced team collaboration, and reusable data transformations. They ensure clear data lineage and enable scalable data preparation solutions, empowering your organization to handle increasingly complex data scenarios with efficiency and clarity.


Divergence

Divergence enables you to create multiple parallel transformation paths from a single step in your workflow. These paths can be transformed independently and later recombined, enabling complex data preparation scenarios such as self-joins.

Creating divergent paths

To initiate a Divergence, in your workflow:

1. Select the step where you want to create divergence.
2. Choose the **+** icon that appears.
3. Configure the new branch that appears.
4. Apply your desired transformations to each path.
5. Use Join or Append steps to recombine paths into a single output.



Key features

- Creates up to five divergent paths from a single step.

- Applies different transformations to each path.
- Recombines paths using Join or Append steps.
- Previews changes in each path independently.

Best practices

- Use divergence for implementing self-joins.
- Create data copies for parallel transformations.
- Plan your recombination strategy (Join or Append).
- Maintain clear path naming for better workflow visibility.

Composite Datasets

Composite Datasets enable you to build upon existing datasets, creating hierarchical data transformation structures that can be shared and reused across your organization. Quick Sight supports up to 10 levels of composite datasets in both SPICE and Direct Query modes.

Creating a composite dataset

To create a composite dataset, in your workflow:

1. Select the Input step when creating a new dataset.
2. Choose **Dataset** as your source under **Add Data**.
3. Select an existing dataset to build upon.
4. Apply additional transformations as needed.
5. Save as a new dataset.

Key features

- Builds hierarchical data transformation structures.
- Supports for up to 10 levels of dataset nesting.
- Compatible with both SPICE and Direct Query.
- Maintains clear data lineage.
- Enables team-specific transformations.

This feature enhances collaboration across different teams. For example,

Role	Action	Output
Global Analyst	Creates dataset with global business logic	Dataset A
Americas Analyst	Uses Dataset A, adds regional logic	Dataset B
US-West Analyst	Uses Dataset B, adds local logic	Dataset C

This hierarchical approach promotes consistent business logic across your organization by assigning clear ownership of transformation layers. It creates a traceable data lineage while supporting up to 10 levels of dataset nesting, enabling controlled and systematic data transformation management.

Best practices

- Establish clear ownership for each transformation layer.
- Document dataset relationships and dependencies.
- Plan hierarchy depth based on business needs.
- Maintain consistent naming conventions.
- Review and update upstream datasets carefully.

SPICE-only features

Amazon Quick Sight's SPICE (Super-fast, Parallel, In-memory Calculation Engine) enables certain computationally intensive data preparation features. These transformations are materialized in SPICE for optimal performance, rather than being executed at query time.

SPICE-only features

Steps	Other capabilities
• Append • Aggregate	• Divergence

Steps	Other capabilities
- Pivot - Unpivot	

Features available in both SPICE and DirectQuery

Steps	Other capabilities
- Input - Add Calculated Columns - Change Data Type - Rename Columns - Select Columns - Filter - Join	Composite Datasets

Best practices

- Use SPICE for workflows requiring SPICE-only features.
- Choose SPICE to optimize performance for complex transformations and large datasets.
- Consider DirectQuery for real-time data needs when SPICE-only features are not required.

Switching between data preparation experiences

Legacy data preparation experience refers to the previous data preparation interface in Amazon Quick Sight that existed before October 2025. The new data preparation experience is the enhanced visual interface that shows step-by-step transformation sequences. Legacy datasets are those created before the new data preparation experience, while new datasets are those created after October 2025.

When creating a new dataset, Quick Sight automatically directs you to the new data preparation experience. This visual interface offers enhanced capabilities and improved usability for data transformation tasks.

Opt-out option

Before saving and publishing a dataset, you have the option to switch back to the legacy data preparation experience, if preferred. This flexibility allows teams to transition at their own pace while becoming familiar with the new interface.

Important

If a dataset is saved and published in the new experience, there will not be an option to go back to the legacy experience. This is by design, as the new experience has significant new features which are not supported in the legacy experience. Hence directly converting datasets from one experience to another is not supported. You will need to create a new dataset to switch to the legacy experience.

Transition workflow

Once a dataset is saved in either the new or legacy experience, the transformations cannot be directly converted from one experience to another. However, if a published dataset version exists, you can use version control to go to the previous version which might be in the legacy experience.

Legacy datasets will continue to be accessible for viewing and editing exclusively through the legacy interface. This maintains compatibility with previously established workflows.

Before fully transitioning, take time to familiarize yourself with the new data preparation experience. When working with legacy datasets, consider creating a new version using the new experience for future modifications. Use version control to maintain access to legacy versions of datasets if needed. Document any changes in workflow when transitioning from legacy to new experience to ensure team alignment.

Features not supported in the new data preparation experience

While the new data preparation experience offers enhanced capabilities, some features from the legacy experience are not yet supported. This section outlines these features and provides guidance for handling affected workflows.

When using unsupported data sources, Amazon Quick Sight automatically defaults to the legacy experience. For other unsupported features, select **Switch to legacy experience** in the top right corner of the data preparation page. Rules Datasets created in the legacy experience remain compatible with both legacy and new experience datasets.

 Note

The new data preparation experience now supports the following features:

- Google Sheets
- Incremental refresh
- JODA date formats
- Column folders
- Column descriptions
- Geospatial data type
- SPICE and Direct Query parent dataset selection from the legacy experience (you can now use a legacy SPICE or Direct Query dataset as a parent for a new experience dataset).

 Setting the geospatial type

You set the geospatial type on the **Output** tab, not through **Change data type** on the **Transform** tab. Geospatial types (such as latitude and longitude) are not part of the regular data types. Because Amazon Quick Sight does not apply further transformations to these columns, it places them on the **Output** tab.

To set the geospatial type of a column:

1. On the data preparation page, choose the **Output** tab at the top (next to **Transform**).
2. In the **Schema** pane on the left, choose the column you want to set (for example, **City**).
3. Open the **Geospatial type** dropdown and choose a value, such as **Country**, **State**, **County**, **City**, **Postcode**, **Latitude**, or **Longitude**. To clear the geospatial type, choose **Unspecified**.

Unsupported data sources

The following data sources are currently available only in the legacy experience.

Data Source	Details
Salesforce	Automatically defaults to legacy experience
S3 Analytics	S3 data sources are supported

Other unsupported features

The following features are currently available only in the legacy experience.

Feature Category	Unsupported features
Dataset Management	Dataset parameters
Data Types	ELF/CLF formats , Zip/GZip files in S3
Configuration Options	"Start from row" in file upload settings
Calculated Fields	UnMaterialized Calculated Columns (UMCC)
Geospatial	Geo Hierarchy

Future development

Amazon Quick Sight plans to implement these features in the new data preparation experience in the future. This approach ensures that the initial launch for the new data preparation experience prioritizes:

Enhanced capabilities

- Visual transformation workflows
- Improved process transparency
- Advanced preparation techniques through Divergence
- Powerful new features like Append, Aggregate, and Pivot

Flexible adoption

Users can choose between experiences before publishing datasets, ensuring uninterrupted workflows while teams transition at their own pace. This approach allows immediate access to new capabilities while maintaining support for specialized requirements through the legacy experience.

Data preparation limits

Amazon Quick Sight's data preparation experience is designed to handle enterprise-scale datasets while maintaining optimal performance. The following limits ensure reliable functionality.

Dataset size limits (SPICE)

- **Output size:** Up to 2TB or 2 billion rows
- **Total input size:** Combined input sources cannot exceed 2TB
- **Secondary tables size:** Combined size is limited to 20GB

Note

Primary tables are those with maximum size in a workflow; all others are secondary.

Workflow structure limits

- **Maximum steps:** Up to 256 transformation steps per workflow
- **Source tables:** Maximum 32 import steps per workflow
- **Output columns:** Up to 2048 columns at any step in the workflow and final output table with 2000 columns
- **Divergent paths:** Maximum 5 paths from a single step (SPICE only, not applicable for DirectQuery)
- **Dataset as a source:** Up to 10 levels for both SPICE and DirectQuery

These limits are designed to balance flexibility with performance, enabling complex data transformations while ensuring optimal analysis capabilities.

Ingestion behavior changes

The new data preparation experience introduces an important change in how data quality issues are handled during SPICE ingestion. This change significantly impacts data completeness and transparency in your datasets.

In the legacy experience, when encountering data type inconsistencies (such as incorrect date formats or [similar issues](#)), the entire row containing problematic cells is skipped during ingestion. This approach results in fewer rows in the final dataset, potentially obscuring data quality issues.

The new experience takes a more granular approach to data inconsistencies. When encountering problematic cells, only the inconsistent values are converted to null values while retaining the entire row. This preservation ensures that related data in other columns remains accessible for analysis.

Impact on dataset quality

Datasets created in the new experience will typically contain more rows than their legacy counterparts when the source data contains inconsistencies. This enhanced approach offers several benefits:

- Improved data completeness by retaining all rows
- Greater transparency in identifying data quality issues
- Better visibility of problematic values for remediation
- Preservation of related data in unaffected columns

This change enables analysts to identify and address data quality issues more effectively, rather than having problematic rows silently omitted from the dataset.

Frequently asked questions

1. When do users need to switch from the new to legacy experience?

Users must return to the legacy experience when working with datasets that contain currently [unsupported features](#). Quick Sight is actively working to incorporate these features into the new experience in upcoming releases.

2. Why are datasets grayed out when trying to add them in the new experience? Can datasets be combined between legacy and new experiences?

Currently, parent and child datasets must exist within the same experience environment. You cannot combine datasets across legacy and new experiences because the new experience includes additional features not available in legacy, such as Append functionalities, Pivot capabilities, and Divergence.

Using parent datasets from the legacy experience

To use parent datasets from the legacy experience, you can switch back to that environment. Simply navigate to the data preparation page and choose **Switch back to legacy experience** in the top right corner. Once there, you can create your child datasets as needed.

Future development

We are planning to implement functionality that will allow users to upgrade legacy datasets to the new experience. This upgraded pathway will enable the use of legacy parent datasets within the new experience.

3. Why is Quick Sight launching the new data preparation experience before achieving full feature parity with the legacy experience?

The new data preparation experience was developed through extensive customer collaboration to address real-world analytics challenges. The initial launch prioritizes:

Enhanced capabilities

- Visual transformation workflows
- Improved process transparency
- Advanced preparation techniques through Divergence
- Powerful new features like Append, Aggregate, and Pivot

Flexible adoption

Users can choose between experiences before publishing datasets, ensuring uninterrupted workflows while teams transition at their own pace. This approach allows immediate access to new capabilities while maintaining support for specialized requirements through the legacy experience.

4. Will features currently available only in the legacy experience be added to the new experience?

Yes. Quick Sight is actively working to incorporate legacy features into the new experience.

5. How do API changes affect existing dataset creation scripts?

Quick Sight maintains backwards compatibility while introducing new capabilities:

- Existing Scripts: Legacy API scripts will continue to function, creating datasets in the legacy experience
- API Naming: Current API names remain unchanged
- New Functionality: Additional API formats support the new experience's enhanced capabilities
- Documentation: Complete API specifications for the new experience are available in our API reference

6. Can datasets be converted between experiences after publication?

- Future Migration Path: Quick Sight will add a feature in the future to easily migrate legacy datasets to the new experience.
- One-Way Process: Converting datasets from the new experience to legacy format isn't supported due to advanced feature dependencies

Dataset Enrichment

Dataset Enrichment is a capability in Amazon Quick Sight that enables dataset authors to add rich semantic metadata to their datasets. By providing descriptions, custom instructions, and structured metadata, you ensure that both human consumers and AI-powered agents understand what a dataset represents and how to use it.

Dataset Enrichment overview

Dataset Enrichment enables authors and author pros to annotate datasets with semantic context at both the dataset level and the column level. This metadata connects raw data with business context. It serves two audiences:

- **Dataset Consumers (other Authors, Reader Pros)** – Gain better business context about what each dataset contains, its purpose, and appropriate use cases.

- **AI Agents** – Receive richer contextual information to generate more accurate queries and interpretations when answering questions through Dataset Q&A.

Dataset Enrichment components

Dataset-level enrichment

Important

Do not add sensitive information to the **Dataset Description** or **Custom Instructions** fields. This information is visible to all dataset viewers.

Dataset Description

A business-level summary of what the dataset represents, its scope, and intended use. This description is visible to all dataset consumers in the UI, helping them quickly understand the dataset's purpose. Maximum length: 5,000 characters.

Custom Instructions

Free-form text instructions specifically consumed by AI agents. These instructions guide the AI on how to interpret, query, and reason about the dataset. Maximum length: 5,000 characters.

File Upload

You can upload a single file in YAML, JSON, or TXT format containing catalog-grade semantic metadata exported from third-party tools (for example, Databricks, dbt, or Alation). This enables hundreds of column definitions, business rules, and metric calculations to be ingested in a single upload – eliminating manual column-by-column entry. Maximum length: 50,000 characters.

Column-level enrichment

Folders

Organize columns into logical groupings for easier navigation and understanding.

Column Description

A human-readable description of what each column represents, its valid values, and business meaning. Maximum length: 500 characters.

Additional Notes

Supplementary context for each column, such as data quality considerations, related tables, or common analysis patterns. Maximum length: 2,000 characters.

Benefits of Dataset Enrichment

- **More accurate AI-powered Dataset Q&A** – Richer semantic context helps AI agents generate more precise SQL queries and interpretations, leading to significantly better answers.
- **Better understanding for consumers** – Descriptions and metadata help all users across the organization understand what datasets contain and how to use them correctly.
- **Scale metadata from external catalogs** – File Upload allows authors to bring in rich metadata from third-party catalog tools in a single operation, rather than manually entering definitions column by column.

Permissions and requirements

Authors and author pros with Enterprise licenses can enrich any dataset they own or manage.

Accessing Dataset Enrichment

To access Dataset Enrichment, complete the following steps.

1. Save your dataset in the data preparation experience.
2. Choose the **Output** tab.
3. Enter the **Dataset Description** and **Custom Instructions**, or upload a semantic metadata file.

Writing effective custom instructions

Custom Instructions are the most impactful component of Dataset Enrichment. They directly guide AI agents on how to interpret and query a dataset. The following are examples of effective and ineffective custom instructions.

Good custom instructions

Example 1 – Revenue Dataset

This dataset contains net revenue after returns and discounts, calculated

on an accrual basis. Revenue is recognized at the point of sale for retail transactions and upon delivery confirmation for B2B orders. All figures are in USD. The 'revenue' column specifically excludes taxes, shipping fees, and promotional credits. For year-over-year comparisons, use the 'fiscal_year' field rather than 'calendar_year' as our fiscal year runs April-March.

Why it's effective:

- Clarifies ambiguous terms (net vs. gross revenue)
- Defines calculation methodology
- Specifies currency and exclusions
- Provides guidance on how to use specific fields correctly

Example 2 – Customer Dataset

Customer status definitions: 'Active' = purchased within last 12 months; 'Dormant' = 12-24 months since last purchase; 'Churned' = 24+ months inactive. The 'customer_segment' field uses RFM analysis (Recency, Frequency, Monetary). 'Lifetime_value' is calculated as total historical spend, not predictive LTV. When analyzing customer counts, always filter out 'is_test_account = true' to exclude internal test data.

Why it's effective:

- Defines business logic and thresholds
- Explains acronyms and methodologies
- Warns about data quality considerations
- Guides proper filtering for accurate analysis

Ineffective custom instructions

Example – Customer Dataset

Contains customer information including names, addresses, purchase history, and other details. Use this for customer analysis.

Why it's ineffective:

- Describes what is already obvious from column names
- Provides no business context or definitions
- Offers no guidance on data quality, calculations, or proper usage
- Does not help the AI distinguish between similar concepts

Key principles for writing good custom instructions

- **Clarify ambiguities** – Define terms that can have multiple interpretations.
- **Explain business logic** – Document calculations, thresholds, and categorizations.
- **Provide context** – Include units, time periods, currencies, and scope.
- **Guide usage** – Explain which fields to use for specific analyses.
- **Warn about edge cases** – Note data quality issues, test records, or special cases.
- **Be specific** – Use concrete examples and precise language.

Two approaches to semantic enrichment

Manual UI-based annotation

Dataset authors directly add dataset and column descriptions and custom instructions through the Quick Sight interface. Quick Sight displays descriptions prominently in the UI, helping all users understand dataset content, column definitions, and appropriate use cases.

File upload from external catalogs

Dataset authors can export semantic metadata from external catalogs and attach a file per dataset in YAML, JSON, or TXT format through the API or UI. While this information is used by AI models rather than displayed in the UI, it enables catalog-grade metadata at scale.

The consumption layer: Dataset Q&A

Dataset Q&A is the consumption layer that uses Dataset Enrichment metadata. It enables users to ask open-ended, natural language questions directly against the datasets they have access to – without needing pre-built dashboards or manually configured topics.

The AI agent uses enriched context in the following ways:

- **Asset discovery** – The agent uses dataset descriptions and semantic metadata to identify the right dataset for the user's question.

- **Text-to-SQL generation** – Custom instructions, column descriptions, and uploaded metadata guide the AI in generating more accurate SQL queries.
- **Governed responses** – All responses respect Row-Level Security (RLS) and Column-Level Security (CLS) rules.

Without enrichment, the AI agent only has column names and data types to work with – which are often ambiguous. With enrichment, the agent receives the full business context needed to:

- Disambiguate similar fields and concepts
- Apply correct calculations and filters
- Understand business-specific thresholds and categorizations
- Exclude test data and handle edge cases appropriately

After you add semantic context to a dataset, users can reference the dataset in Q&A and query it through chat. The AI agent consumes the added metadata to deliver more accurate responses.

Dataset Q&A for Professional subscribers

Dataset Q&A enables users to ask natural language questions directly against a dataset and receive AI-generated answers grounded in the underlying data. Previously, Dataset Q&A was available only to Author personas (Enterprise license holders) who could access datasets to build dashboards and analyses.

Starting June 1, 2026, Dataset View permissions are available for Professional subscribers (Reader Pro). Professional subscribers can now see shared datasets, view the schema (dataset name, description, column names, and column descriptions), and use Chat to ask natural language questions against the data.

Note

This change does not grant Professional subscribers any authoring capabilities. Professional subscribers cannot create analyses, edit datasets, or modify schema. Access is strictly View and Q&A for shared datasets.

Access for already shared datasets

Professional subscribers automatically gain Dataset View access if they meet either of the following criteria:

- They are a member of an Amazon Quick Sight group where a dataset has been shared.
- They are a member of a folder that contains shared datasets.

No action is required from authors or admins – existing sharing configurations automatically extend Dataset View access to Professional subscribers who are already part of these groups or folders.

To share datasets with Professional subscribers

To give Professional subscribers Dataset View access to additional datasets, use one of the following methods:

- Add the Professional subscriber to an Amazon Quick Sight group, then share the dataset with that group. The Professional subscriber automatically receives View access to the dataset.
- Place the dataset in a folder where the Professional subscriber has membership. The Professional subscriber automatically inherits View access through the folder.

In both cases, Professional subscribers receive View-only access. They can see the dataset schema and use Dataset Q&A but cannot edit or create analyses from the dataset.

What Professional subscribers can do with shared datasets

The following list summarizes the permissions that Professional subscribers have for shared datasets.

- **Can** – View the Dataset Details page (see schema including dataset name, description, column names, and column descriptions).
- **Can** – Chat against the dataset using Q&A (ask natural language questions and receive AI-generated answers).
- **Cannot** – Create analyses from the dataset.
- **Cannot** – Edit datasets (no transformations, no schema changes).

Summary

Dataset Enrichment adds semantic metadata to datasets for AI-powered analysis. By investing a few minutes in adding descriptions, custom instructions, and metadata files, dataset authors can improve the accuracy of AI-powered Q&A while making their datasets more understandable and accessible to every consumer across the organization.

Describing data

Using Amazon Quick Sight, you can add information, or *metadata*, about the columns (fields) in your datasets. By adding metadata, you make the dataset self-explanatory and easier to reuse. Doing this can help data curators and their customers know where the data came from and what it means. It's a way of communicating to the people who use your dataset or combine it with other datasets to build dashboards. Metadata is especially important for information that is shared between organizations.

After you add metadata to a dataset, field descriptions become available to anyone who is using the dataset. A column description appears when someone who is actively browsing the **Fields** list pauses on a field name. Column descriptions are visible to people who are editing a dataset or an analysis, but not to someone who is viewing a dashboard. Descriptions aren't formatted. You are able to enter line feeds and formatting marks and these are preserved by the editor. However, the description tooltip that displays is only able to show words, numbers, and symbols—but not formatting.

To edit a description to a column or field

1. From the Quick homepage, choose **Data** at left.
2. In the **Data** tab, choose the dataset that you want to work on.
3. On the dataset details page that opens, choose **Edit dataset** at upper right.
4. On the dataset page that opens, choose a column in the table preview at bottom or in the field list at left.
5. To add or change the description, do one of the following:
 - At the bottom of the screen, open the settings for the field from the pencil icon next to the field's name.
 - In the field list, open the settings for the field from the menu next to the field's name. Then choose **Edit name & description** from the context menu.

6. Add or change the description for the field.

To delete an existing description, delete all the text in the Description box.

7. (Optional) For **Name**, if you want to change the name of the field, you can enter a new one here.
8. Choose **Apply** to save your changes. Choose cancel to exit.

Choosing file upload settings

If you are using a file data source, confirm the upload settings, and correct them if necessary.

Important

If it's necessary to change upload settings, make these changes before you make any other changes to the dataset. Changing upload settings causes Amazon Quick Sight to reimport the file. This process overwrites any changes you have made so far.

Changing text file upload settings

Text file upload settings include the file header indicator, file format, text delimiter, text qualifier, and start row. If you are working with an Amazon S3 data source, the upload settings you select are applied to all files you choose to use in this dataset.

Use the following procedure to change text file upload settings.

1. On the data preparation page, open the **Upload Settings** pane by choosing the expand icon.
2. In **File format**, choose the file format type.
3. If you chose the **custom separated (CUSTOM)** format, specify the separating character in **Delimiter**.
4. If the file doesn't contain a header row, deselect the **Files include headers** check box.
5. If you want to start from a row other than the first row, specify the row number in **Start from row**. If the **Files include headers** check box is selected, the new starting row is treated as the header row. If the **Files include headers** check box is not selected, the new starting row is treated as the first data row.
6. In **Text qualifier**, choose the text qualifier, either single quotes (') or double quotes (").

Changing Microsoft Excel file upload settings

Microsoft Excel file upload settings include the range header indicator and whole worksheet selector.

Use the following procedure to change Microsoft Excel file upload settings.

1. On the data preparation page, open the **Upload Settings** pane by choosing the expand icon.
2. Leave **Upload whole sheet** selected.
3. If the file doesn't contain a header row, deselect the **Range contains headers** check box.

Data Preparation Experience (Legacy)

Topics

- [Adding calculations](#)
- [Joining data](#)
- [Preparing data fields for analysis in Amazon Quick Sight](#)
- [Filtering data in Amazon Quick Sight](#)
- [Previewing tables in a dataset](#)

Adding calculations

Create calculated fields to transform your data by using one or more of the following:

- [Operators](#)
- [Functions](#)
- Fields that contain data
- Other calculated fields

You can add calculated fields to a dataset during data preparation or from the analysis page. When you add a calculated field to a dataset during data preparation, it's available to all analyses that use that dataset. When you add a calculated field to a dataset in an analysis, it's available only in that analysis. For more information about adding calculated fields, see the following topics.

Topics

- [Adding calculated fields](#)
- [Order of evaluation in Amazon Quick Sight](#)
- [Using level-aware calculations in Quick Sight](#)
- [Calculated field function and operator reference for Amazon Quick](#)

Adding calculated fields

Create calculated fields to transform your data by using one or more of the following:

- [Operators](#)
- [Functions](#)
- Aggregate functions (you can only add these to an analysis)
- Fields that contain data
- Other calculated fields

You can add calculated fields to a dataset during data preparation or from the analysis page. When you add a calculated field to a dataset during data preparation, it's available to all analyses that use that dataset. When you add a calculated field to a dataset in an analysis, it's available only in that analysis.

Analyses support both single-row operations and aggregate operations. Single-row operations are those that supply a (potentially) different result for every row. Aggregate operations supply results that are always the same for entire sets of rows. For example, if you use a simple string function with no conditions, it changes every row. If you use an aggregate function, it applies to all the rows in a group. If you ask for the total sales amount for the US, the same number applies to the entire set. If you ask for data on a particular state, the total sales amount changes to reflect your new grouping. It still provides one result for the entire set.

By creating the aggregated calculated field within the analysis, you can then drill down into the data. The value of that aggregated field is recalculated appropriately for each level. This type of aggregation isn't possible during dataset preparation.

For example, let's say that you want to figure out the percentage of profit for each country, region, and state. You can add a calculated field to your analysis, $(\text{sum}(\text{salesAmount} - \text{cost})) / \text{sum}(\text{salesAmount})$. This field is then calculated for each country, region, and state, at the time your analyst drills down into the geography.

Topics

- [Adding calculated fields to an analysis](#)
- [Adding calculated fields to a dataset](#)
- [Handling decimal values in calculated fields](#)

Adding calculated fields to an analysis

When you add a dataset to an analysis, every calculated field that exists in the dataset is added to the analysis. You can add additional calculated fields at the analysis level to create calculated fields that are available only in that analysis.

To add a calculated field to an analysis

1. Open the [Quick console](#).
2. Open the analysis that you want to change.
3. In the **Data** pane, choose **Add** at top left, and then choose **+ CALCULATED FIELD**.
 - a. In the calculations editor that opens, do the following:
 - b. Enter a name for the calculated field.
 - c. Enter a formula using fields from your dataset, functions, and operators.
4. When finished, choose **Save**.

For more information about how to create formulas using the available functions in Quick Sight, see [Calculated field function and operator reference for Amazon Quick](#).

Adding calculated fields to a dataset

Amazon Quick Sight authors can generate calculated fields during the data preparation phase of a dataset's creation. When you create a calculated field for a dataset, the field becomes a new column in the dataset. All analyses that use the dataset inherit the dataset's calculated fields.

If the calculated field operates at the row level and the dataset is stored in SPICE, Quick Sight computes and materializes the result in SPICE. If the calculated field relies on an aggregation function, Quick Sight retains the formula and performs the calculation when the analysis is generated. This type of calculated field is called an unmaterialized calculated field.

To add or edit a calculated field for a dataset

1. Open the dataset that you want to work with. For more information, see [Editing datasets](#).
2. On the data prep page, do one of the following:
 - To create a new field, choose **Add calculated field** at left.
 - To edit an existing calculated field, choose it from **Calculated fields** at left, then choose **Edit** from the context (right-click) menu.
3. In the calculation editor, enter a descriptive name for **Add title** to name the new calculated field. This name appears in the field list in the dataset, so it should look similar to the other fields. For this example, we name the field `Total Sales This Year`.
4. (Optional) Add a comment, for example to explain what the expression does, by enclosing text in slashes and asterisks.

```
/* Calculates sales per year for this year*/
```

5. Identify the metrics, functions, and other items to use. For this example, we need to identify the following:
 - The metric to use
 - Functions: `ifelse` and `datediff`

We want to build a statement like "If the sale happened during this year, show the total sales, and otherwise show 0."

To add the `ifelse` function, open the **Functions** list. Choose **All** to close the list of all functions. Now you should see the function groups: **Aggregate**, **Conditional**, **Date**, and so on.

Choose **Conditional**, and then double-click on `ifelse` to add it to the workspace.

```
ifelse()
```

6. Place your cursor inside the parenthesis in the workspace, and add three blank lines.

```
ifelse(  
  
  

```



```
)
```

7. With your cursor on the first blank line, find the `dateDiff` function. It's listed for **Functions** under **Dates**. You can also find it by entering **date** for **Search functions**. The `dateDiff` function returns all functions that have *date* as part of their name. It doesn't return all functions listed under **Dates**; for example, the `now` function is missing from the search results.

Double-click on `dateDiff` to add it to the first blank line of the `ifelse` statement.

```
ifelse(  
  dateDiff()  
  
)
```

Add the parameters that `dateDiff` uses. Place your cursor inside the `dateDiff` parentheses to begin to add `date1`, `date2`, and `period`:

1. For `date1`: The first parameter is the field that has the date in it. Find it under **Fields**, and add it to the workspace by double-clicking it or entering its name.
2. For `date2`, add a comma, then choose `truncDate()` for **Functions**. Inside its parenthesis, add period and date, like this: `truncDate( "YYYY", now() )`
3. For `period`: Add a comma after `date2` and enter `YYYY`. This is the period for the year. To see a list of all the supported periods, find `dateDiff` in the **Functions** list, and open the documentation by choosing **Learn more**. If you're already viewing the documentation, as you are now, see [dateDiff](#).

Add a few spaces for readability, if you like. Your expression should look like the following.

```
ifelse(  
  dateDiff( {Date}, truncDate( "YYYY", now() ) , "YYYY" )  
  
)
```

8. Specify the return value. For our example, the first parameter in `ifelse` needs to return a value of `TRUE` or `FALSE`. Because we want the current year, and we're comparing it to this year, we specify that the `dateDiff` statement should return `0`. The `if` part of the `ifelse`

evaluates as true for rows where there is no difference between the year of the sale and the current year.

```
dateDiff( {Date}, truncDate( "YYYY", now() ) , "YYYY" ) = 0
```

To create a field for TotalSales for last year, you can change 0 to 1.

Another way to do the same thing is to use addDateTime instead of truncDate. Then for each previous year, you change the first parameter for addDateTime to represent each year. For this, you use -1 for last year, -2 for the year before that, and so on. If you use addDateTime, you leave the dateDiff function = 0 for each year.

```
dateDiff( {Discharge Date}, addDateTime(-1, "YYYY", now() ) , "YYYY" ) = 0 /*
Last year */
```

9. Move your cursor to the first blank line, just under dateDiff. Add a comma.

For the then part of the ifelse statement, we need to choose the measure (metric) that contains the sales amount, TotalSales.

To choose a field, open the **Fields** list and double-click a field to add it to the screen. Or you can enter the name. Add curly braces { } around names that contain spaces. It's likely that your metric has a different name. You can know which field is a metric by the number sign in front of it (#).

Your expression should look like the following now.

```
ifelse(
  dateDiff( {Date}, truncDate( "YYYY", now() ) , "YYYY" ) = 0
  ,{TotalSales}
)
```

10. Add an else clause. The ifelse function doesn't require one, but we want to add it. For reporting purposes, you usually don't want to have any null values, because sometimes rows with nulls are omitted.

We set the else part of the ifelse to 0. The result is that this field is 0 for rows that contain sales from previous years.

To do this, on the blank line add a comma and then a 0. If you added the comment at the beginning, your finished `ifelse` expression should look like the following.

```
/* Calculates sales per year for this year*/
ifelse(
  dateDiff( {Date}, truncDate( "YYYY", now() ) ,"YYYY" ) = 0
  ,{TotalSales}
  ,0
)
```

11. Save your work by choosing **Save** at upper right.

If there are errors in your expression, the editor displays an error message at the bottom. Check your expression for a red squiggly line, then hover your cursor over that line to see what the error message is. Common errors include missing punctuation, missing parameters, misspellings, and invalid data types.

To avoid making any changes, choose **Cancel**.

To add a parameter value to a calculated field

1. You can reference parameters in calculated fields. By adding the parameter to your expression, you add the current value of that parameter.
2. To add a parameter, open the **Parameters** list, and select the parameter whose value you want to include.
3. (Optional) To manually add a parameter to the expression, type the name of the parameter. Then enclosed it in curly braces {}, and prefix it with a \$, for example `${parameterName}`.

You can change the data type of any field in your dataset, including the types of calculated fields. You can only choose data types that match the data that's in the field.

To change the data type of a calculated field

- For **Calculated fields** (at left), choose the field that you want to change, then choose **Change data type** from the context (right-click) menu.

Unlike the other fields in the dataset, calculated fields can't be disabled. Instead, delete them.

To delete a calculated field

- For **Calculated fields** (at left), choose the field that you want to change, then choose **Delete** from the context (right-click) menu.

Handling decimal values in calculated fields

When your dataset uses Direct Query mode, the calculation of the decimal data type is determined by the behavior of the source engine that the dataset originates from. In some particular cases, Quick Sight applies special handlings to determine the output calculation's data type.

When your dataset uses SPICE query mode and a calculated field is materialized, the data type of the result is contingent on the specific function operators and the data type of the input. The tables below show the expected behavior for some numeric calculated fields.

Unary operators

The following table shows which data type is output based on the operator you use and the data type of the value that you input. For example, if you input an integer to an abs calculation, the output value's data type is integer.

Operator	Input type	Output type
abs	Decimal-fixed	Decimal-fixed
	Int	Int
	Decimal-float	Decimal-float
ceil	Decimal-fixed	Int
	Int	Int
	Decimal-float	Int
exp	Decimal-fixed	Decimal-float
	Int	Decimal-float
	Decimal-float	Decimal-float

Operator	Input type	Output type
floor	Decimal-fixed	Int
	Int	Int
	Decimal-float	Int
ln	Decimal-fixed	Decimal-float
	Int	Decimal-float
	Decimal-float	Decimal-float
log	Decimal-fixed	Decimal-float
	Int	Decimal-float
	Decimal-float	Decimal-float
round	Decimal-fixed	Decimal-fixed
	Int	Decimal-fixed
	Decimal-float	Decimal-fixed
sqrt	Decimal-fixed	Decimal-float
	Int	Decimal-float
	Decimal-float	Decimal-float

Binary operators

The following tables show which data type is output based on the data types of the two values that you input. For example, for an arithmetic operator, if you provide two integer data types, the result of the calculation output as an integer.

For basic operators (+, -, *):

	Integer	Decimal-fixed	Decimal-float
Integer	Integer	Decimal-fixed	Decimal-float
Decimal-fixed	Decimal-fixed	Decimal-fixed	Decimal-float
Decimal-float	Decimal-float	Decimal-float	Decimal-float

For division operators (/):

	Integer	Decimal-fixed	Decimal-float
Integer	Decimal-float	Decimal-float	Decimal-float
Decimal-fixed	Decimal-float	Decimal-fixed	Decimal-float
Decimal-float	Decimal-float	Decimal-float	Decimal-float

For exponential and mod operators (^, %):


	Integer	Decimal-fixed	Decimal-float
Integer	Decimal-float	Decimal-float	Decimal-float
Decimal-fixed	Decimal-float	Decimal-float	Decimal-float
Decimal-float	Decimal-float	Decimal-float	Decimal-float

Order of evaluation in Amazon Quick Sight

When you open or update an analysis, before displaying it Amazon Quick Sight evaluates everything that is configured in the analysis in a specific sequence. Amazon Quick Sight translates the configuration into a query that a database engine can run. The query returns the data in a similar way whether you connect to a database, a software as a service (SaaS) source, or the Amazon Quick Sight analytics engine ([SPICE](#)).

If you understand the order that the configuration is evaluated in, you know the sequence that dictates when a specific filter or calculation is applied to your data.

The following illustration shows the order of evaluation. The column on the left shows the order of evaluation when no level aware calculation window (LAC-W) nor aggregate (LAC-A) function is involved. The second column shows the order of evaluation for analyses that contain calculated fields to compute LAC-W expressions at the prefilter (PRE_FILTER) level. The third column shows the order of evaluation for analyses that contain calculated fields to compute LAC-W expressions at the preaggregate (PRE_AGG) level. The last column shows the order of evaluation for analyses that contain calculated fields to compute LAC-A expressions. Following the illustration, there is a more detailed explanation of the order of evaluation. For more information about level aware calculations, see [Using level-aware calculations in Quick Sight](#).



The following list shows the sequence in which Amazon Quick Sight applies the configuration in your analysis. Anything that's set up in your data set happens outside your analysis, for example calculations at the data set level, filters, and security settings. These all apply to the underlying data. The following list only covers what happens inside the analysis.

1. **LAC-W Prefilter level:** Evaluates the data at the original table cardinality before analysis filters
 - a. **Simple calculations:** Calculations at scalar level without any aggregations or window calculations. For example, `date_metric/60`, `parseDate(date, 'yyyy/MM/dd')`, `ifelse(metric > 0, metric, 0)`, `split(string_column, '|' 0)`.
 - b. **LAC-W function PRE_FILTER:** If any LAC-W PRE_FILTER expression is involved in the visual, Amazon Quick Sight firstly computes the window function at the original table level, before

any filters. If the LAC-W PRE_FILTER expression is used in filters, it is applied at this point. For example, `maxOver(Population, [State, County], PRE_FILTER) > 1000`.

2. **LAC-W PRE_AGG:** Evaluates the data at the original table cardinality before aggregations

- a. **Filters added during analysis:** Filters created for un-aggregated fields in the visuals are applied at this point, which are similar to WHERE clauses. For example, `year > 2020`.
- b. **LAC-W function PRE_AGG:** If any LAC-W PRE_AGG expression is involved in the visual, Amazon Quick Sight computes the window function before any aggregation applied. If the LAC-W PRE_AGG expression is used in filters, it is applied at this point. For example, `maxOver(Population, [State, County], PRE_AGG) > 1000`.
- c. **Top/bottom N filters:** Filters that are configured on dimensions to display top/bottom N items.

3. **LAC-A level:** Evaluate aggregations at customized level, before visual aggregations

- a. **Custom-level aggregations:** If any LAC-A expression is involved in the visual, it is calculated at this point. Based on the table after the filters mentioned above, Amazon QuickSight computes the aggregation, grouped by the dimensions that are specified in the calculated fields. For example, `max(Sales, [Region])`.

4. **Visual level:** Evaluates aggregations at visual level, and post-aggregation table calculations, with the remaining configurations applied in the visuals

- a. **Visual-level aggregations:** Visual aggregations should always be applied except for tabular tables (where dimension is empty). With this setting, aggregations based on the fields in the field wells are calculated, grouped by the dimensions that put into the visuals. If any filter is built on top of aggregations, it is applied at this point, similar to HAVING clauses. For example, `min(distance) > 100`.
- b. **Table calculations:** If there is any post-aggregation table calculation (it should take aggregated expression as operand) referenced in the visual, it is calculated at this point. Amazon Quick Sight performs window calculations after visual aggregations. Similarly, filters built on such calculations are applied.
- c. **Other category calculations:** This type of calculation only exists in line/bar/pie/donut charts. For more information, see [Display limits](#).
- d. **Totals and subtotals:** Totals and Subtotals are calculated in donut charts (only totals), tables (only totals) and pivot tables, if requested.

Using level-aware calculations in Quick Sight

Applies to: Enterprise Edition and Standard Edition

With *Level-aware calculations* (LAC) you can specify the level of granularity that you want to compute window functions or aggregate functions. There are two types of LAC functions: level-aware calculation - aggregate (LAC-A) functions, and level-aware calculation - window (LAC-W) functions.

Topics

- [Level-aware calculation - aggregate \(LAC-A\) functions](#)
- [Level-aware calculation - window \(LAC-W\) functions](#)

Level-aware calculation - aggregate (LAC-A) functions

With LAC-A functions, you can specify at what level to group the computation. By adding one argument into an existing aggregate function, such as `sum()` , `max()` , `count()`, you can define any group-by level that you want for the aggregation. The level added can be any dimension independent of the dimensions added to the visual. For example:

```
sum(measure,[group_field_A])
```

To use LAC-A functions, type them directly in the calculation editor by adding the intended aggregation levels as the second argument between brackets. Following is an example of an aggregate function and a LAC-A function, for comparison.

- Aggregate function: `sum({sales})`
- LAC-A function: `sum({sales}, [{Country},{Product}])`

The LAC-A results are computed with the specified level in the brackets [], can be used as operand of an aggregate function. The group-by level of the aggregate function is visual level, with **Group by** fields added to the field well of the visual.

In addition to creating a static LAC group key in the bracket [], you can make it dynamically adapted to visual group-by fields, by putting a parameter `$visualDimensions` in the

bracket. This is a system-provided parameter, in contrast to user-defined parameter. The `[$visualDimensions]` parameter represents the fields added to the **Group by** field well in current visual. The following examples show how to dynamically add group keys to the visual dimensions or remove group keys from visual dimensions

- LAC-A with dynamic-added group key : `sum({sales}, [{visualDimensions}, {Country}, {Products}])`

It calculates, before the visual level aggregation is calculated, the sum of sales, grouping by country, products, and any other fields in the **Group by** field well .

- LAC-A with dynamic-removed group key : `sum({sales}, [{visualDimensions}, !{Country}, !{Products}])`

It calculates, before visual level aggregation is calculated, the sum of sales, grouping by the fields in the visual's **Group by** field well, except country and product.

You can specify added group key or removed group key in on LAC expression, but not both.

LAC-A functions are supported for the following aggregate functions:

- [avg](#)
- [count](#)
- [distinct_count](#)
- [max](#)
- [median](#)
- [min](#)
- [percentile](#)
- [percentileCont](#)
- [percentileDisc \(percentile\)](#)
- [stdev](#)
- [stdevp](#)
- [sum](#)
- [var](#)
- [varp](#)

LAC-A examples

You can do the following with LAC-A functions:

- Run calculations that are independent of the levels in the visual. For example, if you have the following calculation, the sales numbers are aggregated only at the country level, but not across other dimensions (Region or Product) in the visual.

```
sum({Sales}, [{Country}])
```

- Run calculations for the dimensions that are not in the visual. For example, if you have the following function, you can calculate the average total country sales by region.

```
sum({Sales}, [{Country}])
```

Though Country is not included in the visual, the LAC-A function first aggregates the sales at the Country level and then the visual level calculation generates the average number for each region. If the LAC-A function isn't used to specify the level, the average sales are calculated at the lowest granular level (the base level of the dataset) for each region (showing in the sales column).

- Use LAC-A combined with other aggregate functions and LAC-W functions. There are two ways you can nest LAC-A functions with other functions.
 - You can write a nested syntax when you create a calculation. For example, the LAC-A function can be nested with a LAC-W function to calculate the total sales by country of each product's average price:

```
sum(avgOver({Sales}, [{Product}], PRE_AGG), [{Country}])
```

- When adding a LAC-A function into a visual, the calculation can be further nested with visual-level aggregate functions that you selected in the fields well. For more information about changing the aggregation of fields in the visual, see [Changing or adding aggregation to a field by using a field well](#).

LAC-A limitations

The following limitations apply to LAC-A functions:

- LAC-A functions are supported for all additive and non-additive aggregate functions, such as `sum()`, `count()`, and `percentile()`. LAC-A functions are not supported for conditional aggregate functions that end with "if", such as `sumif()` and `countif()`, nor for period aggregate functions that start with "periodToDate", such as `periodToDateSum()` and `periodToDateMax()`.
- Row-level and column-level totals are not currently supported for LAC-A functions in tables and pivot tables. When you add row-level or column-level totals to the chart, the total number will show as blank. Other non-LAC dimensions are not impacted.
- Nested LAC-A functions are not currently supported. A limited capability of LAC-A functions nested with regular aggregate functions and LAC-W functions are supported.

For example, the following functions are valid:

- `Aggregation(LAC-A())`. For example: `max(sum({sales}, [{country}]))`
- `LAC-A(LAC-W())`. For example: `sum(sumOver({Sales}, [{Product}], PRE_AGG), [{Country}])`

The following functions are not valid:

- `LAC-A(Aggregation())`. For example: `sum(max({sales}), [{country}])`
- `LAC-A(LAC-A())`. For example: `sum(max({sales}, [{country}]), [category])`
- `LAC-W(LAC-A())`. For example: `sumOver(sum({Sales}, [{Product}]), [{Country}], PRE_AGG)`

Level-aware calculation - window (LAC-W) functions

With LAC-W functions, you can specify the window or partition to compute the calculation. LAC-W functions are a group of window functions, such as `sumover()`, `(maxover)`, `denseRank`, that you can run at the prefilter or preaggregate level. For example: `sumOver(measure, [partition_field_A], pre_agg)`.

LAC-W functions used to be called level aware aggregations (LAA).

LAC-W functions help you to answer the following types of questions:

- How many of my customers made only 1 purchase order? Or 10? Or 50? We want the visual to use the count as a dimension rather than a metric in the visual.
- What are the total sales per market segment for customers whose lifetime spend is greater than \$100,000? The visual should only show the market segment and the total sales for each.

- How much is the contribution of each industry to the entire company's profit (percent of total)? We want to be able to filter the visual to show some of the industries, and how they contribute to the total sales for the displayed industries. However, we also want to see each industry's percent of total sales for the entire company (including industries that are filtered out).
- What are the total sales in each category as compared to the industry average? The industry average should include all of the categories, even after filtering.
- How are my customers grouped into cumulative spending ranges? We want to use the grouping as a dimension rather than a metric.

For more complex questions, you can inject a calculation or filter before Quick Sight gets to a specific point in its evaluation of your settings. To directly influence your results, you add a calculation level keyword to a table calculation. For more information on how Quick Sight evaluates queries, see [Order of evaluation in Amazon Quick Sight](#).

The following calculation levels are supported for LAC-W functions:

- **PRE_FILTER** – Before applying filters from the analysis, Quick Sight evaluates prefilter calculations. Then it applies any filters that are configured on these prefilter calculations.
- **PRE_AGG** – Before computing display-level aggregations, Quick Sight performs preaggregate calculations. Then it applies any filters that are configured on these preaggregate calculations. This work happens before applying top and bottom *N* filters.

You can use the PRE_FILTER or PRE_AGG keyword as a parameter in the following table calculation functions. When you specify a calculation level, you use an unaggregated measure in the function. For example, you can use `countOver({ORDER ID}, [{Customer ID}], PRE_AGG)`. By using PRE_AGG, you specify that the `countOver` executes at the preaggregate level.

- [avgOver](#)
- [countOver](#)
- [denseRank](#)
- [distinctCountOver](#)
- [minOver](#)
- [maxOver](#)
- [percentileRank](#)
- [rank](#)

- [stdevOver](#)
- [stdevpOver](#)
- [sumOver](#)
- [varOver](#)
- [varpOver](#)

By default, the first parameter for each function must be an aggregated measure. If you use either PRE_FILTER or PRE_AGG, you use a nonaggregated measure for the first parameter.

For LAC-W functions, the visual aggregation defaults to MIN to eliminate duplicates. To change the aggregation, open the field's context (right-click) menu, and then choose a different aggregation.

For examples of when and how to use LAC-W functions in real life scenarios, see the following post in the AWS Big Data Blog: [Create advanced insights using Level Aware Aggregations in Amazon QuickSight](#).

Calculated field function and operator reference for Amazon Quick

You can add calculated fields to a dataset during data preparation or from the analysis page. When you add a calculated field to a dataset during data preparation, it's available to all analyses that use that dataset. When you add a calculated field to a dataset in an analysis, it's available only in that analysis.

You can create calculated fields to transform your data by using the following functions and operators.

Topics

- [Operators](#)
- [Functions by category](#)
- [Functions](#)
- [Aggregate functions](#)
- [Table calculation functions](#)

Operators

You can use the following operators in calculated fields. Quick uses the standard order of operations: parentheses, exponents, multiplication, division, addition, subtraction (PEMDAS). Equal (=) and not equal (<>) comparisons are case-sensitive.

- Addition (+)
- Subtraction (-)
- Multiplication (*)
- Division (/)
- Modulo (%) – See also `mod( )` in the following list.
- Power (^) – See also `exp( )` in the following list.
- Equal (=)
- Not equal (<>)
- Greater than (>)
- Greater than or equal to (>=)
- Less than (<)
- Less than or equal to (<=)
- AND
- OR
- NOT

Amazon Quick supports applying the following mathematical functions to an expression.

- [Mod](#)(*number*, *divisor*) – Finds the remainder after dividing a number by a divisor.
- [Log](#)(*expression*) – Returns the base 10 logarithm of a given expression.
- [Ln](#)(*expression*) – Returns the natural logarithm of a given expression.
- [Abs](#)(*expression*) – Returns the absolute value of a given expression.
- [Sqrt](#)(*expression*) – Returns the square root of a given expression.
- [Exp](#)(*expression*) – Returns the base of natural log *e* raised to the power of a given expression.

To make lengthy calculations easier to read, you can use parenthesis to clarify groupings and precedence in calculations. In the following statement, you don't need parentheses. The multiplication statement is processed first, and then the result is added to five, returning a value of 26. However, parentheses make the statement easier to read and thus maintain.

```
5 + (7 * 3)
```

Because parenthesis are first in the order of operations, you can change the order in which other operators are applied. For example, in the following statement the addition statement is processed first, and then the result is multiplied by three, returning a value of 36.

```
(5 + 7) * 3
```

Example: Arithmetic operators

The following example uses multiple arithmetic operators to determine a sales total after discount.

```
(Quantity * Amount) - Discount
```

Example: (/) Division

The following example uses division to divide 3 by 2. A value of 1.5 is returned. Amazon Quick uses floating point divisions.

```
3/2
```

Example: (=) equal

Using = performs a case-sensitive comparison of values. Rows where the comparison is TRUE are included in the result set.

In the following example, rows where the Region field is **South** are included in the results. If the Region is **south**, these rows are excluded.

```
Region = 'South'
```

In the following example, the comparison evaluates to FALSE.

```
Region = 'south'
```

The following example shows a comparison that converts `Region` to all uppercase (**SOUTH**), and compares it to **SOUTH**. This returns rows where the region is **south**, **South**, or **SOUTH**.

```
toUpper(Region) = 'SOUTH'
```

Example: (<>)

The not equal symbol `<>` means *less than or greater than*.

So, if we say **x<>1**, then we are saying *if x is less than 1 OR if x is greater than 1*. Both `<` and `>` are evaluated together. In other words, *if x is any value except 1*. Or, *x is not equal to 1*.

Note

Use `<>`, not `!=`.

The following example compares `Status Code` to a numeric value. This returns rows where the `Status Code` is not equal to **1**.

```
statusCode <> 1
```

The following example compares multiple `statusCode` values. In this case, active records have `activeFlag = 1`. This example returns rows where one of the following applies:

- For active records, show rows where the status isn't 1 or 2
- For inactive records, show rows where the status is 99 or -1

```
( activeFlag = 1 AND (statusCode <> 1 AND statusCode <> 2) )  
OR  
( activeFlag = 0 AND (statusCode= 99 OR statusCode= -1) )
```

Example: (^)

The power symbol `^` means *to the power of*. You can use the power operator with any numeric field, with any valid exponent.

The following example is a simple expression of 2 to the power of 4 or $(2 * 2 * 2 * 2)$. This returns a value of 16.

```
2^4
```

The following example computes the square root of the revenue field.

```
revenue^0.5
```

Example: AND, OR, and NOT

The following example uses AND, OR, and NOT to compare multiple expressions. It does so using conditional operators to tag top customers NOT in Washington or Oregon with a special promotion, who made more than 10 orders. If no values are returned, the value 'n/a' is used.

```
ifelse(( NOT (State = 'WA' OR State = 'OR')) AND Orders > 10), 'Special Promotion  
XYZ', 'n/a')
```

Example: Creating comparison lists like "in" or "not in"

This example uses operators to create a comparison to find values that exist, or don't exist, in a specified list of values.

The following example compares promoCode a specified list of values. This example returns rows where the promoCode is in the list **(1, 2, 3)**.

```
promoCode    = 1  
OR promoCode = 2  
OR promoCode = 3
```

The following example compares promoCode a specified list of values. This example returns rows where the promoCode is NOT in the list **(1, 2, 3)**.

```
NOT(promoCode = 1  
OR promoCode = 2  
OR promoCode = 3  
)
```

Another way to express this is to provide a list where the promoCode is not equal to any items in the list.

```
promoCode    <> 1
```

```
AND promoCode <> 2  
AND promoCode <> 3
```

Example: Creating a "between" comparison

This example uses comparison operators to create a comparison showing values that exist between one value and another.

The following example examines `OrderDate` and returns rows where the `OrderDate` is between the first day and last day of 2016. In this case, we want the first and last day included, so we use "or equal to" on the comparison operators.

```
OrderDate >= "1/1/2016" AND OrderDate <= "12/31/2016"
```

Functions by category

In this section, you can find a list of the functions available in Amazon Quick, sorted by category.

Topics

- [Aggregate functions](#)
- [Conditional functions](#)
- [Date functions](#)
- [Numeric functions](#)
- [Mathematical functions](#)
- [String functions](#)
- [Table calculations](#)

Aggregate functions

The aggregate functions for calculated fields in Amazon Quick include the following. These are only available during analysis and visualization. Each of these functions returns values grouped by the chosen dimension or dimensions. For each aggregation, there is also a conditional aggregation. These perform the same type of aggregation, based on a condition.

- [avg](#) averages the set of numbers in the specified measure, grouped by the chosen dimension or dimensions.
- [avgIf](#) calculates the average based on a conditional statement.

- [count](#) calculates the number of values in a dimension or measure, grouped by the chosen dimension or dimensions.
- [countIf](#) calculates the count based on a conditional statement.
- [distinct_count](#) calculates the number of distinct values in a dimension or measure, grouped by the chosen dimension or dimensions.
- [distinct_countIf](#) calculates the distinct count based on a conditional statement.
- [max](#) returns the maximum value of the specified measure, grouped by the chosen dimension or dimensions.
- [maxIf](#) calculates the maximum based on a conditional statement.
- [median](#) returns the median value of the specified measure, grouped by the chosen dimension or dimensions.
- [medianIf](#) calculates the median based on a conditional statement.
- [min](#) returns the minimum value of the specified measure, grouped by the chosen dimension or dimensions.
- [minIf](#) calculates the minimum based on a conditional statement.
- [percentile](#) (alias of `percentileDisc`) computes the *n*th percentile of the specified measure, grouped by the chosen dimension or dimensions.
- [percentileCont](#) calculates the *n*th percentile based on a continuous distribution of the numbers of the specified measure, grouped by the chosen dimension or dimensions.
- [percentileDisc \(percentile\)](#) calculates the *n*th percentile based on the actual numbers of the specified measure, grouped by the chosen dimension or dimensions.
- [periodToDateAvg](#) averages the set of numbers in the specified measure for a given time granularity (for instance, a quarter) up to a point in time.
- [periodToDateCount](#) calculates the number of values in a dimension or measure for a given time granularity (for instance, Quarter) up to a point in time including duplicates.
- [periodToDateMax](#) returns the maximum value of the specified measure for a given time granularity (for instance, a quarter) up to a point in time.
- [periodToDateMedian](#) returns the median value of the specified measure for a given time granularity (for instance, a quarter) up to a point in time.
- [periodToDateMin](#) returns the minimum value of the specified measure or date for a given time granularity (for instance, a quarter) up to a point in time.
- [periodToDatePercentile](#) calculates the percentile based on the actual numbers in measure for a given time granularity (for instance, a quarter) up to a point in time.

- [periodToDatePercentileCont](#) calculates percentile based on a continuous distribution of the numbers in the measure for a given time granularity (for instance, a quarter) up to a point in time.
- [periodToDateStDev](#) calculates the standard deviation of the set of numbers in the specified measure for a given time granularity (for instance, a quarter) up to a point in time based on a sample.
- [periodToDateStDevP](#) calculates the population standard deviation of the set of numbers in the specified measure for a given time granularity (for instance, a quarter) up to a point in time based on a sample.
- [periodToDateSum](#) adds the set of numbers in the specified measure for a given time granularity (for instance, a quarter) up to a point in time.
- [periodToDateVar](#) calculates the sample variance of the set of numbers in the specified measure for a given time granularity (for instance, a quarter) up to a point in time.
- [periodToDateVarP](#) calculates the population variance of the set of numbers in the specified measure for a given time granularity (for instance, a quarter) up to a point in time.
- [stdev](#)) calculates the standard deviation of the set of numbers in the specified measure, grouped by the chosen dimension or dimensions, based on a sample.
- [stdevIf](#) calculates the sample standard deviation based on a conditional statement.
- [stdevp](#) calculates the standard deviation of the set of numbers in the specified measure, grouped by the chosen dimension or dimensions, based on a biased population.
- [stdevpIf](#) calculates the population deviation based on a conditional statement.
- [var](#)) calculates the variance of the set of numbers in the specified measure, grouped by the chosen dimension or dimensions, based on a sample.
- [varIf](#) calculates the sample variance based on a conditional statement.
- [varp](#)) calculates the variance of the set of numbers in the specified measure, grouped by the chosen dimension or dimensions, based on a biased population.
- [varpIf](#) calculates the population variance based on a conditional statement.
- [sum](#)) adds the set of numbers in the specified measure, grouped by the chosen dimension or dimensions.
- [sumIf](#)) calculates the sum based on a conditional statement.

Conditional functions

The conditional functions for calculated fields in Amazon Quick include the following:

- [Coalesce](#) returns the value of the first argument that is not null.
- [Ifelse](#) evaluates a set of *if, then* expression pairings, and returns the value of the *then* argument for the first *if* argument that evaluates to true.
- [in](#) evaluates an expression to see if it is in a given list of values.
- [isNotNull](#) evaluates an expression to see if it is not null.
- [isNull](#) evaluates an expression to see if it is null. If the expression is null, `isNull` returns true, and otherwise it returns false.
- [notIn](#) evaluates an expression to see if it is not in a given list of values.
- [nullIf](#) compares two expressions. If they are equal, the function returns null. If they are not equal, the function returns the first expression.
- [switch](#) returns an expression that matches the first label equal to the condition expression.

Date functions

The date functions for calculated fields in Amazon Quick include the following:

- [addDateTime](#) adds or subtracts a unit of time to the date or time provided.
- [addWorkDays](#) adds or subtracts the given number of work days to the date or time provided.
- [dateDiff](#) returns the difference in days between two date fields.
- [epochDate](#) converts an epoch date into a standard date.
- [Extract](#) returns a specified portion of a date value.
- [formatDate](#) formats a date using a pattern you specify.
- [isWorkDay](#) returns TRUE if a given date-time value is a work or business day.
- [netWorkDays](#) returns the number of working days between the provided two date values.
- [Now](#) returns the current date and time, using either settings for a database, or UTC for file and Salesforce.
- [truncDate](#) returns a date value that represents a specified portion of a date.

Numeric functions

The numeric functions for calculated fields in Amazon Quick include the following:

- [Ceil](#) rounds a decimal value to the next highest integer.
- [decimalToInt](#) converts a decimal value to an integer.

- [Floor](#) decrements a decimal value to the next lowest integer.
- [intToDecimal](#) converts an integer value to a decimal.
- [Round](#) rounds a decimal value to the closest integer or, if scale is specified, to the closest decimal place.

Mathematical functions

The mathematical functions for calculated fields in Amazon Quick include the following:

- [Mod](#)(*number*, *divisor*) – Finds the remainder after dividing a number by a divisor.
- [Log](#)(*expression*) – Returns the base 10 logarithm of a given expression.
- [Ln](#)(*expression*) – Returns the natural logarithm of a given expression.
- [Abs](#)(*expression*) – Returns the absolute value of a given expression.
- [Sqrt](#)(*expression*) – Returns the square root of a given expression.
- [Exp](#)(*expression*) – Returns the base of natural log e raised to the power of a given expression.

String functions

The string (text) functions for calculated fields in Amazon Quick include the following:

- [Concat](#) concatenates two or more strings.
- [contains](#) checks if an expression contains a substring.
- [endsWith](#) checks if the expression ends with the substring specified.
- [Left](#) returns the specified number of leftmost characters from a string.
- [Locate](#) locates a substring within another string, and returns the number of characters before the substring.
- [Ltrim](#) removes preceding blank space from a string.
- [parseDate](#) parses a string to determine if it contains a date value, and returns the date if found.
- [parseDecimal](#) parses a string to determine if it contains a decimal value.
- [parseInt](#) parses a string to determine if it contains an integer value.
- [parseJson](#) parses values from a native JSON or from a JSON object in a text field.
- [Replace](#) replaces part of a string with a new string.
- [Right](#) returns the specified number of rightmost characters from a string.

- [Rtrim](#) removes following blank space from a string.
- [Split](#) splits a string into an array of substrings, based on a delimiter that you choose, and returns the item specified by the position.
- [startsWith](#) checks if the expression starts with the substring specified.
- [Strlen](#) returns the number of characters in a string.
- [Substring](#) returns the specified number of characters in a string, starting at the specified location.
- [toLowerCase](#) formats a string in all lowercase.
- [toString](#) formats the input expression as a string.
- [toUpperCase](#) formats a string in all uppercase.
- [trim](#) removes both preceding and following blank space from a string.

Table calculations

Table calculations form a group of functions that provide context in an analysis. They provide support for enriched aggregated analysis. By using these calculations, you can address common business scenarios such as calculating percentage of total, running sum, difference, common baseline, and rank.

When you are analyzing data in a specific visual, you can apply table calculations to the current set of data to discover how dimensions influence measures or each other. Visualized data is your result set based on your current dataset, with all the filters, field selections, and customizations applied. To see exactly what this result set is, you can export your visual to a file. A table calculation function performs operations on the data to reveal relationships between fields.

Lookup-based functions

- [difference](#) calculates the difference between a measure based on one set of partitions and sorts, and a measure based on another.
- [lag](#) calculates the lag (previous) value for a measure.
- [lead](#) calculates the lead (following) value for a measure.
- [percentDifference](#) calculates the percentage difference between the current value and a comparison value.

Over functions

- [avgOver](#) calculates the average of a measure over one or more dimensions.
- [countOver](#) calculates the count of a field over one or more dimensions.
- [distinctCountOver](#) calculates the distinct count of the operand partitioned by the specified attributes at a specified level.
- [maxOver](#) calculates the maximum of a measure over one or more dimensions.
- [minOver](#) the minimum of a measure over one or more dimensions.
- [percentileOver](#) (alias of `percentileDiscOver`) calculates the n th percentile of a measure partitioned by a list of dimensions.
- [percentileContOver](#) calculates the n th percentile based on a continuous distribution of the numbers of a measure partitioned by a list of dimensions.
- [percentileDiscOver](#) calculates the n th percentile based on the actual numbers of a measure partitioned by a list of dimensions.
- [percentOfTotal](#) calculates the percentage that a measure contributes to the total.
- [periodOverPeriodDifference](#) calculates the difference of a measure over two different time periods as specified by period granularity and offset.
- [periodOverPeriodLastValue](#) calculates the last (previous) value of a measure from a previous time period as specified by period granularity and offset.
- [periodOverPeriodPercentDifference](#) calculates the percent difference of a measure over two different time periods as specified by period granularity and offset.
- [periodToDateAvgOverTime](#) calculates the average of a measure for a given time granularity (for instance, a quarter) up to a point in time.
- [periodToDateCountOverTime](#) calculates the count of a dimension or measure for a given time granularity (for instance, a quarter) up to a point in time.
- [periodToDateMaxOverTime](#) calculates the maximum of a measure or date for a given time granularity (for instance, a quarter) up to a point in time.
- [periodToDateMinOverTime](#) calculates the minimum of a measure or date for a given time granularity (for instance, a quarter) up to a point in time.
- [periodToDateSumOverTime](#) calculates the sum of a measure for a given time granularity (for instance, a quarter) up to a point in time.
- [sumOver](#) calculates the sum of a measure over one or more dimensions.
- [stdevOver](#) calculates the standard deviation of the specified measure, partitioned by the chosen attribute or attributes, based on a sample.

- [stdevpOver](#) calculates the standard deviation of the specified measure, partitioned by the chosen attribute or attributes, based on a biased population.
- [varOver](#) calculates the variance of the specified measure, partitioned by the chosen attribute or attributes, based on a sample.
- [varpOver](#) calculates the variance of the specified measure, partitioned by the chosen attribute or attributes, based on a biased population.

Ranking functions

- [rank](#) calculates the rank of a measure or a dimension.
- [denseRank](#) calculates the rank of a measure or a dimension, ignoring duplicates.
- [percentileRank](#) calculates the rank of a measure or a dimension, based on percentile.

Running functions

- [runningAvg](#) calculates a running average for a measure.
- [runningCount](#) calculates a running count for a measure.
- [runningMax](#) calculates a running maximum for a measure.
- [runningMin](#) calculates a running minimum for a measure.
- [runningSum](#) calculates a running sum for a measure.

Window functions

- [firstValue](#) calculates the first value of the aggregated measure or dimension partitioned and sorted by specified attributes.
- [lastValue](#) calculates the last value of the aggregated measure or dimension partitioned and sorted by specified attributes.
- [windowAvg](#) calculates the average of the aggregated measure in a custom window that is partitioned and sorted by specified attributes.
- [windowCount](#) calculates the count of the aggregated measure in a custom window that is partitioned and sorted by specified attributes.
- [windowMax](#) calculates the maximum of the aggregated measure in a custom window that is partitioned and sorted by specified attributes.

- [windowMin](#) calculates the minimum of the aggregated measure in a custom window that is partitioned and sorted by specified attributes.
- [windowSum](#) calculates the sum of the aggregated measure in a custom window that is partitioned and sorted by specified attributes.

Functions

In this section, you can find a list of functions available in Amazon Quick. To view a list of functions sorted by category, with brief definitions, see [Functions by category](#).

Topics

- [addDateTime](#)
- [addWorkDays](#)
- [Abs](#)
- [Ceil](#)
- [Coalesce](#)
- [Concat](#)
- [contains](#)
- [decimalToInt](#)
- [dateDiff](#)
- [endsWith](#)
- [epochDate](#)
- [Exp](#)
- [Extract](#)
- [Floor](#)
- [formatDate](#)
- [Ifelse](#)
- [in](#)
- [intToDecimal](#)
- [isNotNull](#)
- [isNull](#)
- [isWorkDay](#)

- [Left](#)
- [Locate](#)
- [Log](#)
- [Ln](#)
- [Ltrim](#)
- [Mod](#)
- [netWorkDays](#)
- [Now](#)
- [notIn](#)
- [nullIf](#)
- [parseDate](#)
- [parseDecimal](#)
- [parseInt](#)
- [parseJson](#)
- [Replace](#)
- [Right](#)
- [Round](#)
- [Rtrim](#)
- [Split](#)
- [Sqrt](#)
- [startsWith](#)
- [Strlen](#)
- [Substring](#)
- [switch](#)
- [toLower](#)
- [toString](#)
- [toUpper](#)
- [trim](#)
- [truncDate](#)

addDateTime

`addDateTime` adds or subtracts a unit of time from a datetime value. For example, `addDateTime(2, 'YYYY', parseDate('02-JUL-2018', 'dd-MMM-yyyy'))` returns 02-JUL-2020. You can use this function to perform date math on your date and time data.

Syntax

```
addDateTime(amount, period, datetime)
```

Arguments

amount

A positive or negative integer value that represents the amount of time that you want to add or subtract from the provided datetime field.

period

A positive or negative value that represents the amount of time that you want to add or subtract from the provided datetime field. Valid periods are as follows:

- YYYY: This returns the year portion of the date.
- Q: This returns the quarter that the date belongs to (1–4).
- MM: This returns the month portion of the date.
- DD: This returns the day portion of the date.
- WK: This returns the week portion of the date. The week starts on Sunday in Amazon Quick.
- HH: This returns the hour portion of the date.
- MI: This returns the minute portion of the date.
- SS: This returns the second portion of the date.
- MS: This returns the millisecond portion of the date.

datetime

The date or time that you want to perform date math on.

Return type

Datetime

Example

Let's say you have a field called `purchase_date` that has the following values.

```
2018 May 13 13:24
2017 Jan 31 23:06
2016 Dec 28 06:45
```

Using the following calculations, `addDateTime` modifies the values as shown following.

```
addDateTime(-2, 'YYYY', purchaseDate)
```

```
2016 May 13 13:24
2015 Jan 31 23:06
2014 Dec 28 06:45
```

```
addDateTime(4, 'DD', purchaseDate)
```

```
2018 May 17 13:24
2017 Feb 4 23:06
2017 Jan 1 06:45
```

```
addDateTime(20, 'MI', purchaseDate)
```

```
2018 May 13 13:44
2017 Jan 31 23:26
2016 Dec 28 07:05
```

addWorkDays

`addWorkDays` Adds or subtracts a designated number of work days to a given date value. The function returns a date for a work day, that falls a designated work days after or before a given input date value.

Syntax

```
addWorkDays(initDate, numWorkDays)
```

Arguments

initDate

A valid non-NULL date that acts as the start date for the calculation.

- **Dataset field** – Any date field from the dataset that you are adding this function to.
- **Date function** – Any date output from another date function, for example `parseDate`, `epochDate`, `addDateTime`., and so on.

Example

```
addWorkDays(epochDate(1659484800), numWorkDays)
```

- **Calculated fields** – Any Quick calculated field that returns a date value.

Example

```
calcFieldStartDate = addDateTime(10, "DD", startDate)  
addWorkDays(calcFieldStartDate, numWorkDays)
```

- **Parameters** – Any Quick datetime parameter.

Example

```
addWorkDays($paramStartDate, numWorkDays)
```

- Any combination of the above stated argument values.

numWorkDays

A non-NULL integer that acts as the end date for the calculation.

- **Literal** – An integer literal directly typed in the expression editor.

Example

- **Dataset field** – Any date field from the dataset

Example

- **Scalar function or calculation** – Any scalar Quick function that returns an integer output from another, for example `decimalToInt`, `abs`, and so on.

Example

```
addWorkDays(initDate, decimalToInt(sqrt (abs(numWorkDays))) ) )
```

- **Calculated field** – Any Quick calculated field that returns a date value.

Example

```
someOtherIntegerCalcField = (num_days * 2) + 12  
addWorkDays(initDate, someOtherIntegerCalcField)
```

- **Parameter** – Any Quick datetime parameter.

Example

```
addWorkDays(initDate, $param_numWorkDays)
```

- Any combination of the above stated argument values.

Return type

Integer

Output values

Expected output values include:

- Positive integer (when `start_date < end_date`)
- Negative integer (when `start_date > end_date`)
- NULL when one or both of the arguments get a null value from the `dataset field`.

Input errors

Disallowed argument values cause errors, as shown in the following examples.

- Using a literal NULL as an argument in the expression is disallowed.

Example

```
addWorkDays(NULL, numWorkDays)
```

Example

Error

At least one of the arguments in this function does not have correct type.
Correct the expression and choose Create again.

- Using a string literal as an argument, or any other data type other than a date, in the expression is disallowed. In the following example, the string "2022-08-10" looks like a date, but it is actually a string. To use it, you would have to use a function that converts to a date data type.

Example

```
addWorkDays("2022-08-10", 10)
```

Example

Error

Expression addWorkDays("2022-08-10", numWorkDays) for function addWorkDays has incorrect argument type addWorkDays(String, Number).
Function syntax expects Date, Integer.

Example

A positive integer as numWorkDays argument will yield a date in the future of the input date. A negative integer as numWorkDays argument will yield a resultant date in the past of the input date. A zero value for the numWorkDays argument yields the same value as input date whether or not it falls on a work day or a weekend.

The addWorkDays function operates at the granularity: DAY. Accuracy cannot be preserved at any granularity which is lower or higher than DAY level.

```
addWorkDays(startDate, endDate)
```

Let's assume there is a field named employmentStartDate with the following values:

```
2022-08-10 2022-08-06 2022-08-07
```

Using the above field and following calculations, `addWorkDays` returns the modified values as shown below:

```
addWorkDays(employmentStartDate, 7)
```

```
2022-08-19
```

```
2022-08-16
```

```
2022-08-16
```

```
addWorkDays(employmentStartDate, -5)
```

```
2022-08-02
```

```
2022-08-01
```

```
2022-08-03
```

```
addWorkDays(employmentStartDate, 0)
```

```
2022-08-10
```

```
2022-08-06
```

```
2022-08-07
```

The following example calculates the total pro-rated bonus to be paid to each employee for 2 years based on how many days each employee has actually worked.

```
last_day_of_work = addWorkDays(employment_start_date, 730)
total_days_worked = netWorkDays(employment_start_date, last_day_of_work)
total_bonus = total_days_worked * bonus_per_day
```

Field wells

Group by

- employee_id
- employment_end_date
- last_day_of_work

Value

- total_days_worked (Sum)
- bonus_per_day (Sum)
- total_bonus (Sum)

netWorkDays addWorkDays

Total pro-rated bonus expenditure

employee...	employment_end_date	last_day_of_work	total_days_worked	bonus_per_day	total_bonus
101	May 6, 2022	Oct 20, 2023	731	55	40,205
102	May 15, 2022	Oct 23, 2023	731	55	40,205
103	Aug 3, 2022	Oct 27, 2023	730	55	40,150
104	Jan 8, 2022	Oct 27, 2023	730	55	40,150
105	Jan 14, 2022	May 21, 2024	731	55	40,205
106	Sep 2, 2022	May 24, 2024	730	55	40,150
107	Aug 9, 2022	May 27, 2024	731	55	40,205
			5,114	385	281,270

Abs

abs returns the absolute value of a given expression.

Syntax

```
abs(expression)
```

Arguments

expression

The expression must be numeric. It can be a field name, a literal value, or another function.

Ceil

ceil rounds a decimal value to the next highest integer. For example, ceil(29.02) returns 30.

Syntax

```
ceil(decimal)
```

Arguments

decimal

A field that uses the decimal data type, a literal value like **17.62**, or a call to another function that outputs a decimal.

Return type

Integer

Example

The following example rounds a decimal field to the next highest integer.

```
ceil(salesAmount)
```

The following are the given field values.

```
20.13  
892.03  
57.54
```

For these field values, the following values are returned.

```
21  
893  
58
```

Coalesce

`coalesce` returns the value of the first argument that is not null. When a non-null value is found, the remaining arguments in the list are not evaluated. If all arguments are null, the result is null. 0-length strings are valid values and are not considered equivalent to null.

Syntax

```
coalesce(expression1, expression2 [, expression3, ...])
```

Arguments

coalesce takes two or more expressions as arguments. All of the expressions must have the same data type or be able to be implicitly cast to the same data type.

expression

The expression can be numeric, datetime, or string. It can be a field name, a literal value, or another function.

Return type

coalesce returns a value of the same data type as the input arguments.

Example

The following example retrieves a customer's billing address if it exists, her street address if there is no billing address, or returns "No address listed" if neither address is available.

```
coalesce(billingAddress, streetAddress, 'No address listed')
```

Concat

concat concatenates two or more strings.

Syntax

```
concat(expression1, expression2 [, expression3 ...])
```

Arguments

concat takes two or more string expressions as arguments.

expression

The expression must be a string. It can be the name of a field that uses the string data type, a literal value like '**12 Main Street**', or a call to another function that outputs a string.

Return type

String

Examples

The following example concatenates three string fields and adds appropriate spacing.

```
concat(salutation, ' ', firstName, ' ', lastName)
```

The following are the given field values.

salutation	firstName	lastName
Ms.	Li	Juan
Dr.	Ana Carolina	Silva
Mr.	Nikhil	Jayashankar

For these field values, the following values are returned.

```
Ms. Li Juan  
Dr. Ana Carolina Silva  
Mr. Nikhil Jayashankar
```

The following example concatenates two string literals.

```
concat('Hello', 'world')
```

The following value is returned.

```
Helloworld
```

contains

`contains` evaluates if the substring that you specify exists within an expression. If the expression contains the substring, `contains` returns true, and otherwise it returns false.

Syntax

```
contains(expression, substring, string-comparison-mode)
```

Arguments

expression

The expression must be a string. It can be the name of a field that uses the string data type, a literal value like **'12 Main Street'**, or a call to another function that outputs a string.

substring

The set of characters to check against the *expression*. The substring can occur one or more times in the *expression*.

string-comparison-mode

(Optional) Specifies the string comparison mode to use:

- CASE_SENSITIVE – String comparisons are case-sensitive.
- CASE_INSENSITIVE – String comparisons are case-insensitive.

This value defaults to CASE_SENSITIVE when blank.

Return type

Boolean

Examples

Default case sensitive example

The following case sensitive example evaluates if `state_nm` contains **New**.

```
contains(state_nm, "New")
```

The following are the given field values.

```
New York  
new york
```

For these field values, the following values are returned.

```
true
```



```
false
```

Case insensitive example

The following case insensitive example evaluates if `state_nm` contains **new**.

```
contains(state_nm, "new", CASE_INSENSITIVE)
```

The following are the given field values.

```
New York  
new york
```

For these field values, the following values are returned.

```
true  
true
```

Example with conditional statements

The `contains` function can be used as the conditional statement within the following If functions: [avgIf](#), [minIf](#), [distinct_countIf](#), [countIf](#), [maxIf](#), [medianIf](#), [stdevIf](#), [stdevpIf](#), [sumIf](#), [varIf](#), and [varpIf](#).

The following example sums `Sales` only if `state_nm` contains **New**.

```
sumIf(Sales, contains(state_nm, "New"))
```

Does NOT contain example

The conditional NOT operator can be used to evaluate if the expression does not contain the specified substring.

```
NOT(contains(state_nm, "New"))
```

Example using numeric values

Numeric values can be used in the expression or substring arguments by applying the `toString` function.

```
contains(state_nm, toString(5) )
```

decimalToInt

`decimalToInt` converts a decimal value to the integer data type by stripping off the decimal point and any numbers after it. `decimalToInt` does not round up. For example, `decimalToInt(29.99)` returns 29.

Syntax

```
decimalToInt(decimal)
```

Arguments

decimal

A field that uses the decimal data type, a literal value like **17.62**, or a call to another function that outputs a decimal.

Return type

Integer

Example

The following example converts a decimal field to an integer.

```
decimalToInt(salesAmount)
```

The following are the given field values.

```
20.13  
892.03  
57.54
```

For these field values, the following values are returned.

```
20  
892  
57
```

dateDiff

dateDiff returns the difference in days between two date fields. If you include a value for the period, dateDiff returns the difference in the period interval, rather than in days.

Syntax

```
dateDiff(date1, date2, [period])
```

Arguments

dateDiff takes two dates as arguments. Specifying a period is optional.

date 1

The first date in the comparison. A date field or a call to another function that outputs a date.

date 2

The second date in the comparison. A date field or a call to another function that outputs a date.

period

The period of difference that you want returned, enclosed in quotes. Valid periods are as follows:

- YYYY: This returns the year portion of the date.
- Q: This returns the date of the first day of the quarter that the date belongs to.
- MM: This returns the month portion of the date.
- DD: This returns the day portion of the date.
- WK: This returns the week portion of the date. The week starts on Sunday in Amazon Quick.
- HH: This returns the hour portion of the date.
- MI: This returns the minute portion of the date.
- SS: This returns the second portion of the date.
- MS: This returns the millisecond portion of the date.

Return type

Integer

Example

The following example returns the difference between two dates.

```
dateDiff(orderDate, shipDate, "MM")
```

The following are the given field values.

orderDate	shipdate
01/01/18	03/05/18
09/13/17	10/20/17

For these field values, the following values are returned.

```
2  
1
```

endsWith

`endsWith` evaluates if the expression ends with a substring that you specify. If the expression ends with the substring, `endsWith` returns true, and otherwise it returns false.

Syntax

```
endsWith(expression, substring, string-comparison-mode)
```

Arguments

expression

The expression must be a string. It can be the name of a field that uses the string data type, a literal value like **'12 Main Street'**, or a call to another function that outputs a string.

substring

The set of characters to check against the *expression*. The substring can occur one or more times in the *expression*.

string-comparison-mode

(Optional) Specifies the string comparison mode to use:

- **CASE_SENSITIVE** – String comparisons are case-sensitive.

- `CASE_INSENSITIVE` – String comparisons are case-insensitive.

This value defaults to `CASE_SENSITIVE` when blank.

Return type

Boolean

Examples

Default case sensitive example

The following case sensitive example evaluates if `state_nm` endsWith **"York"**.

```
endsWith(state_nm, "York")
```

The following are the given field values.

```
New York  
new york
```

For these field values, the following values are returned.

```
true  
false
```

Case insensitive example

The following case insensitive example evaluates if `state_nm` endsWith **"york"**.

```
endsWith(state_nm, "york", CASE_INSENSITIVE)
```

The following are the given field values.

```
New York  
new york
```

For these field values, the following values are returned.

```
true  
true
```

Example with conditional statements

The `endsWith` function can be used as the conditional statement within the following If functions: [avgIf](#), [minIf](#), [distinct_countIf](#), [countIf](#), [maxIf](#), [medianIf](#), [stdevIf](#), [stdevplf](#), [sumIf](#), [varIf](#), and [varplf](#).

The following example sums Sales only if `state_nm` ends with "York".

```
sumIf(Sales, endsWith(state_nm, "York"))
```

Does NOT contain example

The conditional NOT operator can be used to evaluate if the expression does not start with the specified substring.

```
NOT(endsWith(state_nm, "York"))
```

Example using numeric values

Numeric values can be used in the expression or substring arguments by applying the `toString` function.

```
endsWith(state_nm, toString(5) )
```

epochDate

`epochDate` converts an epoch date into a standard date in the format `yyyy-MM-ddTkk:mm:ss.SSSZ`, using the format pattern syntax specified in [Class DateTimeFormat](#) in the Joda project documentation. An example is `2015-10-15T19:11:51.003Z`.

`epochDate` is supported for use with analyses based on datasets stored in Quick (SPICE).

Syntax

```
epochDate(epochdate)
```

Arguments

epochdate

An epoch date, which is an integer representation of a date as the number of seconds since 00:00:00 UTC on January 1, 1970.

epochdate must be an integer. It can be the name of a field that uses the integer data type, a literal integer value, or a call to another function that outputs an integer. If the integer value is longer than 10 digits, the digits after the 10th place are discarded.

Return type

Date

Example

The following example converts an epoch date to a standard date.

```
epochDate(3100768000)
```

The following value is returned.

```
2068-04-04T12:26:40.000Z
```

Exp

exp returns the base of natural log e raised to the power of a given expression.

Syntax

```
exp(expression)
```

Arguments

expression

The expression must be numeric. It can be a field name, a literal value, or another function.

Extract

extract returns a specified portion of a date value. Requesting a time-related portion of a date that doesn't contain time information returns 0.

Syntax

```
extract(period, date)
```

Arguments

period

The period that you want extracted from the date value. Valid periods are as follows:

- YYYY: This returns the year portion of the date.
- Q: This returns the quarter that the date belongs to (1–4).
- MM: This returns the month portion of the date.
- DD: This returns the day portion of the date.
- WD: This returns the day of the week as an integer, with Sunday as 1.
- HH: This returns the hour portion of the date.
- MI: This returns the minute portion of the date.
- SS: This returns the second portion of the date.
- MS: This returns the millisecond portion of the date.



Note

Extracting milliseconds is not supported in Presto databases below version 0.216.

date

A date field or a call to another function that outputs a date.

Return type

Integer

Example

The following example extracts the day from a date value.

```
extract('DD', orderDate)
```

The following are the given field values.

```
orderDate
=====
01/01/14
```



```
09/13/16
```

For these field values, the following values are returned.

```
01  
13
```

Floor

`floor` decrements a decimal value to the next lowest integer. For example, `floor(29.08)` returns 29.

Syntax

```
floor(decimal)
```

Arguments

decimal

A field that uses the decimal data type, a literal value like **17.62**, or a call to another function that outputs a decimal.

Return type

Integer

Example

The following example decrements a decimal field to the next lowest integer.

```
floor(salesAmount)
```

The following are the given field values.

```
20.13  
892.03  
57.54
```

For these field values, the following values are returned.

```
20
892
57
```

formatDate

`formatDate` formats a date using a pattern you specify. When you are preparing data, you can use `formatDate` to reformat the date. To reformat a date in an analysis, you choose the format option from the context menu on the date field.

Syntax

```
formatDate(date, ['format'])
```

Arguments

date

A date field or a call to another function that outputs a date.

format

(Optional) A string containing the format pattern to apply. This argument accepts the format patterns specified in [Supported date formats](#).

If you don't specify a format, this string defaults to `yyyy-MM-ddTkk:mm:ss.SSS`.

Return type

String

Example

The following example formats a UTC date.

```
formatDate(orderDate, 'dd-MMM-yyyy')
```

The following are the given field values.

```
order date
=====
```

```
2012-12-14T00:00:00.000Z
2013-12-29T00:00:00.000Z
2012-11-15T00:00:00.000Z
```

For these field values, the following values are returned.

```
13 Dec 2012
28 Dec 2013
14 Nov 2012
```

Example

If the date contains single quotes or apostrophes, for example `yyyyMMdd'T'HHmmss`, you can handle this date format by using one of the following methods.

- Enclose the entire date in double quotes, as shown in the following example:

```
formatDate({myDateField}, "yyyyMMdd'T'HHmmss")
```

- Escape the single quotes or apostrophes by adding a backslash (\) to the left of them, as shown in the following example:

```
formatDate({myDateField}, 'yyyyMMdd\'T\'HHmmss')
```

Ifelse

`ifelse` evaluates a set of *if, then* expression pairings, and returns the value of the *then* argument for the first *if* argument that evaluates to true. If none of the *if* arguments evaluate to true, then the value of the *else* argument is returned.

Syntax

```
ifelse(if-expression-1, then-expression-1 [, if-expression-n, then-expression-n
...], else-expression)
```

Arguments

`ifelse` requires one or more *if, then* expression pairings, and requires exactly one expression for the *else* argument.

if-expression

The expression to be evaluated as true or not. It can be a field name like **address1**, a literal value like '**Unknown**', or another function like `toString(salesAmount)`. An example is `isNotNull(FieldName)`.

If you use multiple AND and OR operators in the `if` argument, enclose statements in parentheses to identify processing order. For example, the following `if` argument returns records with a month of 1, 2, or 5 and a year of 2000.

```
ifelse((month = 5 OR month < 3) AND year = 2000, 'yes', 'no')
```

The next `if` argument uses the same operators, but returns records with a month of 5 and any year, or with a month of 1 or 2 and a year of 2000.

```
ifelse(month = 5 OR (month < 3 AND year = 2000), 'yes', 'no')
```

then-expression

The expression to return if its *if* argument is evaluated as true. It can be a field name like **address1**, a literal value like '**Unknown**', or a call to another function. The expression must have the same data type as the other then arguments and the `else` argument.

else-expression

The expression to return if none of the *if* arguments evaluate as true. It can be a field name like **address1**, a literal value like '**Unknown**', or another function like `toString(salesAmount)`. The expression must have the same data type as all of the then arguments.

Return type

`ifelse` returns a value of the same data type as the values in *then-expression*. All data returned *then* and *else* expressions must be of the same data type or be converted to the same data type.

Examples

The following example generates a column of aliases for field `country`.

```
ifelse(country = "United States", "US", country = "China", "CN", country = "India",  
      "IN", "Others")
```

For such use cases evaluating each value in a field against a list of literals, and returns the result corresponding to the first matching value., function `switch` is recommended to simplify your work. The previous example can be rewritten to the following statement using [switch](#):

```
switch(country,"United States","US","China","CN","India","IN","Others")
```

The following example categorizes sales per customer into human-readable levels.

```
ifelse(salesPerCustomer < 1000, "VERY_LOW", salesPerCustomer < 10000, "LOW",
salesPerCustomer < 100000, "MEDIUM", "HIGH")
```

The following example uses AND, OR, and NOT to compare multiple expressions using conditional operators to tag top customers NOT in Washington or Oregon with a special promotion, who made more than 10 orders. If no values are returned, the value 'n/a' is used.

```
ifelse(( NOT (State = 'WA' OR State = 'OR')) AND Orders > 10), 'Special Promotion
XYZ', 'n/a')
```

The following examples use only OR to generate a new column that contains the name of continent that corresponds to each country.

```
ifelse(country = "United States" OR country = "Canada", "North America", country =
"China" OR country = "India" OR country = "Japan", "Asia", "Others")
```

The previous example can be simplified as shown in the next example. The following example uses `ifelse` and [in](#) to create a value in a new column for any row where the tested value is in a literal list. You could use `ifelse` with [notIn](#) as well.

```
ifelse(in(country,["United States", "Canada"]), "North America", in(country,
["China","Japan","India"]), "Asia", "Others")
```

Authors are able to save a literal list in a multivalue parameter and use it in the [in](#) or [notIn](#) functions. The following example is an equivalent of the previous example, except that the literal lists are stored in two multivalue parameters.

```
ifelse(in(country,${NorthAmericaCountryParam}), "North America", in(country,
${AsiaCountryParam}), "Asia", "Others")
```

The following example assigns a group to a sales record based on the sales total. The structure of each if-then phrase mimics the behavior of *between*, a keyword that doesn't currently work in calculated field expressions. For example, the result of the comparison `salesTotal >= 0 AND salesTotal < 500` returns the same values as the SQL comparison `salesTotal between 0 and 499`.

```
ifelse(salesTotal >= 0 AND salesTotal < 500, 'Group 1', salesTotal >= 500 AND
salesTotal < 1000, 'Group 2', 'Group 3')
```

The following example tests for a NULL value by using `coalesce` to return the first non-NULL value. Instead of needing to remember the meaning of a NULL in a date field, you can use a readable description instead. If the disconnect date is NULL, the example returns the suspend date, unless both of those are NULL. Then `coalesce(DiscoDate, SuspendDate, '12/31/2491')` returns `'12/31/2491'`. The return value must match the other data types. This date might seem like an unusual value, but a date in the 25th century reasonably simulates the "end of time," defined as the highest date in a data mart.

```
ifelse ( (coalesce(DiscoDate, SuspendDate, '12/31/2491') = '12/31/2491'), 'Active
subscriber', 'Inactive subscriber')
```

The following shows a more complex example in a more readable format, just to show that you don't need to compress your code all into one long line. This example provides for multiple comparisons of the value a survey result. It handles potential NULL values for this field and categorizes two acceptable ranges. It also labels one range that needs more testing and another that's not valid (out of range). For all remaining values, it applies the `else` condition, and labels the row as needing a retest three years after the date on that row.

```
ifelse
(
  isNull({SurveyResult}), 'Untested',
  {SurveyResult}=1, 'Range 1',
  {SurveyResult}=2, 'Range 2',
  {SurveyResult}=3, 'Need more testing',
  {SurveyResult}=99, 'Out of Range',
  concat
  (
    'Retest by ',
    toString
    (
```

```

        addDateTime(3, "YYYY", {Date}))
    )
)

```

The following example assigns a "manually" created region name to a group of states. It also uses spacing and comments, wrapped in `/* */`, to make it easier to maintain the code.

```

ifelse
(
  /* NE REGION*/
  locate('New York, New Jersey, Connecticut, Vermont, Maine, Rhode Island, New
Hampshire',{State})) > 0,
  'Northeast',

  /* SE REGION*/
  locate('Georgia, Alabama, South Carolina, Louisiana',{State})) > 0,
  'Southeast',

  'Other Region'
)

```

The logic for the region tagging breaks down as follows:

1. We list the states that we want for each region, enclosing each list in quotation marks to make each list a string, as follows:
 - 'New York, New Jersey, Connecticut, Vermont, Maine, Rhode Island, New Hampshire'
 - 'Georgia, Alabama, South Carolina, Louisiana'
 - You can add more sets, or use countries, cities, provinces, or What3Words if you want.
2. We ask if the value for State (for each row) is found in the list, by using the `locate` function to return a nonzero value if the state is found in the list, as follows.

```
locate('New York, New Jersey, Connecticut, Vermont, Maine, Rhode Island, New
Hampshire',{State})
```

and

```
locate('Georgia, Alabama, South Carolina, Louisiana',{State})
```

3. The `locate` function returns a number instead of a `TRUE` or `FALSE`, but `ifelse` requires the `TRUE/FALSE` Boolean value. To get around this, we can compare the result of `locate` to a number. If the state is in the list, the return value is greater than zero.

- a. Ask if the state is present.

```
locate('New York, New Jersey, Connecticut, Vermont, Maine, Rhode Island, New
Hampshire',{State}) > 0
```

- b. If it's present the region, label it as the specific region, in this case a Northeast region.

```
/*The if expression:*/      locate('New York, New Jersey, Connecticut, Vermont,
Maine, Rhode Island, New Hampshire',{State}) > 0,
/*The then expression:*/    'Northeast',
```

4. Because we have states that aren't in a list, and because `ifelse` requires a single `else` expression, we provide `'Other Region'` as the label for the leftover states.

```
/*The if expression:*/      locate('New York, New Jersey, Connecticut, Vermont, Maine,
Rhode Island, New Hampshire',{State}) > 0,
/*The then expression:*/    'Northeast',
/*The else expression:*/    'Other Region'
```

5. We wrap all that in the `ifelse( )` function to get the final version. The following example leaves out the Southeast region states that were in the original. You can add them back in place of the *<insert more regions here>* tag.

If you want to add more regions, you can construct more copies of those two lines and alter the list of states to suit your purpose. You can change the region name to something that suits you, and change the field name from `State` to anything that you need.

```
ifelse
(
/*The if expression:*/      locate('New York, New Jersey, Connecticut, Vermont, Maine,
Rhode Island, New Hampshire',{State}) > 0,
/*The then expression:*/    'Northeast',

/*<insert more regions here>*/

/*The else expression:*/    'Other Region'
)
```


Note

There are other ways to do the initial comparison for the if expression. For example, suppose that you pose the question "What states are not missing from this list?" rather than "Which states are on the list?" If you do, you might phrase it differently. You might compare the locate statement to zero to find values that are missing from the list, and then use the NOT operator to classify them as "not missing," as follows.

```
/*The if expression:*/      NOT (locate('New York, New Jersey, Connecticut,  
Vermont, Maine, Rhode Island, New Hampshire',{State}) = 0),
```

Both versions are correct. The version that you choose should make the most sense to you and your team, so you can maintain it easily. If all the options seem equal, choose the simplest.

in

`in` evaluates if an expression exists within a literal list. If the list contains the expression, `in` returns true, and otherwise it returns false. `in` is case sensitive for string type inputs.

`in` accepts two kinds of literal list, one is manually entered list and the other is a [multivalue parameter](#).

Syntax

Using a manually entered list:

```
in(expression, [literal-1, ...])
```

Using a multivalue parameter:

```
in(expression, $multivalue_parameter)
```

Arguments

expression

The expression to be compared with the elements in literal list. It can be a field name like address, a literal value like '**Unknown**', a single value parameter, or a call to another scalar function—provided this function is not an aggregate function or a table calculation.

literal list

(required) This can be a manually entered list or a multivalue parameter. This argument accepts up to 5,000 elements. However, in a direct query to a third party data source, for example Oracle or Teradata, the restriction can be smaller.

- **manually entered list** – One or more literal values in a list to be compared with the expression. The list should be enclosed in square brackets. All the literals to compare must have the same datatype as the expression.
- **multivalue parameter** – A pre-defined multivalue parameter passed in as a literal list. The multivalue parameter must have the same datatype as the expression.

Return type

Boolean: TRUE/FALSE

Example with a static list

The following example evaluates the `origin_state_name` field for values in a list of string. When comparing string type input, `in` only supports case sensitive comparison.

```
in(origin_state_name,["Georgia", "Ohio", "Texas"])
```

The following are the given field values.

```
"Washington"  
  "ohio"  
  "Texas"
```

For these field values the following values are returned.

```
false
```

```
false
true
```

The third return value is true because only "Texas" is one of the included values.

The following example evaluates the `fl_date` field for values in a list of string. In order to match the type, `toString` is used to cast the date type to string type.

```
in(toString(fl_date),["2015-05-14","2015-05-15","2015-05-16"])
```

Field wells

Group by
fl_date
flightInMidMay

Value
Add measures here

Sheet 1 +

Fl_date and Flightinmidmay

fl_date	flightInMidMay
May 12, 2015	0
May 13, 2015	0
May 15, 2015	1

Literals and NULL values are supported in expression argument to be compared with the literals in list. Both of the following two examples will generate a new column of TRUE values.

```
in("Washington",["Washington","Ohio"])
```

```
in(NULL,[NULL,"Ohio"])
```

Example with multivalue parameter

Let's say an author creates a [multivalue parameter](#) that contains a list of all the state names. Then the author adds a control to allow the reader to select values from the list.

Next, the reader selects three values—"Georgia", "Ohio", and "Texas"—from the parameter's drop down list control. In this case, the following expression is equivalent to the first example, where those three state names are passed as the literal list to be compared with the `original_state_name` field.

```
in (origin_state_name, ${stateName MultivalueParameter})
```

Example with ifelse

`in` can be nested in other functions as a boolean value. One example is that authors can evaluate any expression in a list and return the value they want by using `in` and `ifelse`. The following example evaluates if the `dest_state_name` of a flight is in a particular list of US states and returns different categories of the states based on the comparison.

```
ifelse(in(dest_state_name, ["Washington", "Oregon", "California"]), "WestCoastUSState",  
      "Other US State")
```

Field wells

Group by

dest_state_nm

StateCategory

Value

Add measures here

Sheet 1

Dest_state_nm and Statecategory

dest_state_nm	StateCategory
Alabama	Other US State
Alaska	Other US State
California	WestCoastUSState
Colorado	Other US State
Connecticut	Other US State
Florida	Other US State
Georgia	Other US State
Hawaii	Other US State

intToDecimal

intToDecimal converts an integer value to the decimal data type.

Syntax

```
intToDecimal(integer)
```

Arguments

int

A field that uses the integer data type, a literal value like **14**, or a call to another function that outputs an integer.

Return type

Decimal(Fixed) in the legacy data preparation experience.

Decimal(Float) in the new data preparation experience.

Example

The following example converts an integer field to a decimal.

```
intToDecimal(price)
```

The following are the given field values.

```
20  
892  
57
```

For these field values, the following values are returned.

```
20.0  
892.0  
58.0
```

You can apply formatting inside an analysis, for example to format `price` as currency.

isNotNull

`isNotNull` evaluates an expression to see if it is not null. If the expression is not null, `isNotNull` returns true, and otherwise it returns false.

Syntax

```
isNotNull(expression)
```

Arguments

expression

The expression to be evaluated as null or not. It can be a field name like **address1** or a call to another function that outputs a string.

Return type

Boolean

Example

The following example evaluates the `sales_amount` field for null values.

```
isNotNull(salesAmount)
```

The following are the given field values.

```
20.13  
(null)  
57.54
```

For these field values, the following values are returned.

```
true  
false  
true
```

isNull

`isNull` evaluates an expression to see if it is null. If the expression is null, `isNull` returns true, and otherwise it returns false.

Syntax

```
isNull(expression)
```

Arguments

expression

The expression to be evaluated as null or not. It can be a field name like **address1** or a call to another function that outputs a string.

Return type

Boolean

Example

The following example evaluates the `sales_amount` field for null values.

```
isNull(salesAmount)
```

The following are the given field values.

```
20.13  
(null)  
57.54
```

For these field values, the following values are returned.

```
false  
true  
false
```

The following example tests for a NULL value in an `ifelse` statement, and returns a human-readable value instead.

```
ifelse( isNull({ActiveFlag}) , 'Inactive', 'Active')
```

isWorkDay

`isWorkDay` evaluates a given date-time value to determine if the value is a workday or not.

`isWorkDay` assumes a standard 5-day work week starting from Monday and ending on Friday. Saturday and Sunday are assumed to be weekends. The function always calculates its result at the DAY granularity and is exclusive of the given input date.

Syntax

```
isWorkDay(inputDate)
```

Arguments

inputDate

The date-time value that you want to evaluate. Valid values are as follows:

- **Dataset fields:** Any date field from the dataset that you are adding this function to.
- **Date Functions:** Any date output from another date function, for example, `parseDate`.
- **Calculated fields:** Any Quick calculated field that returns a date value.
- **Parameters:** Any Quick `DateTime` parameter.

Return type

Integer (0 or 1)

Example

The following example determines whether or not the `application_date` field is a work day.

Let's assume that there's a field named `application_date` with the following values:

```
2022-08-10  
2022-08-06  
2022-08-07
```

When you use these fields and add the following calculations, `isWorkDay` returns the below values:

```
isWorkDay({application_date})  
  
1  
0  
0
```

The following example filters employees whose employment ends on a work day and determines whether their employment began on work day or a weekend using conditional formatting:

```
is_start_date_work_day = isWorkDay(employment_start_date)  
is_end_date_work_day = isWorkDay(employment_end_date)
```

Field wells

Group by

employee_id

employment_start_date

employment_end_date

is_start_date_work_day

is_end_date_work_day

Value

Add measures here

netWorkDays

addWorkDays

isWorkDay

+

Filter employees with end date on week day and discern b/w start dates

employee...	employment_start_date	employment_end_date	is_start_date_work_day	is_end_date_work_day
101	Jan 1, 2021	May 6, 2022	1	1
103	Jan 9, 2021	Aug 3, 2022	0	1
105	Aug 3, 2021	Jan 14, 2022	1	1
106	Aug 7, 2021	Sep 2, 2022	0	1
107	Aug 9, 2021	Aug 9, 2022	1	1

Left

left returns the leftmost characters from a string, including spaces. You specify the number of characters to be returned.

Syntax

```
left(expression, limit)
```

Arguments

expression

The expression must be a string. It can be the name of a field that uses the string data type, a literal value like **'12 Main Street'**, or a call to another function that outputs a string.

limit

The number of characters to be returned from *expression*, starting from the first character in the string.

Return type

String

Example

The following example returns the first 3 characters from a string.

```
left('Seattle Store #14', 3)
```

The following value is returned.

```
Sea
```

Locate

locate locates a substring that you specify within another string, and returns the number of characters until the first character in the substring. The function returns 0 if it doesn't find the substring. The function is 1-based.

Syntax

```
locate(expression, substring, start)
```

Arguments

expression

The expression must be a string. It can be the name of a field that uses the string data type, a literal value like **'12 Main Street'**, or a call to another function that outputs a string.

substring

The set of characters in *expression* that you want to locate. The substring can occur one or more times in *expression*.

start

(Optional) If *substring* occurs more than once, use *start* to identify where in the string the function should start looking for the substring. For example, suppose that you want to find the second example of a substring and you think it typically occurs after the first 10 characters. You specify a *start* value of 10. It should start from 1.

Return type

Integer

Examples

The following example returns information about where the first occurrence of the substring 'and' appears in a string.

```
locate('1 and 2 and 3 and 4', 'and')
```

The following value is returned.

```
3
```

The following example returns information about where the first occurrence of the substring 'and' appears in a string after the fourth character.

```
locate('1 and 2 and 3 and 4', 'and', 4)
```

The following value is returned.

```
9
```

Log

log returns the base 10 logarithm of a given expression.

Syntax

```
log(expression)
```

Arguments

expression

The expression must be numeric. It can be a field name, a literal value, or another function.

Ln

Ln returns the natural logarithm of a given expression.

Syntax

```
ln(expression)
```

Arguments

expression

The expression must be numeric. It can be a field name, a literal value, or another function.

Ltrim

Ltrim removes preceding blank space from a string.

Syntax

```
ltrim(expression)
```

Arguments

expression

The expression must be a string. It can be the name of a field that uses the string data type, a literal value like **'12 Main Street'**, or a call to another function that outputs a string.

Return type

String

Example

The following example removes the preceding spaces from a string.

```
ltrim('   Seattle Store #14')
```

The following value is returned.

```
Seattle Store #14
```

Mod

Use the mod function to find the remainder after dividing the number by the divisor. You can use the mod function or the modulo operator (%) interchangeably.

Syntax

```
mod(number, divisor)
```

```
number%divisor
```

Arguments

number

The number is the positive integer that you want to divide and find the remainder for.

divisor

The divisor is the positive integer that you are dividing by. If the divisor is zero, this function returns an error on dividing by 0.

Example

The following examples return the modulo of 17 when dividing by 6. The first example uses the % operator, and the second example uses the mod function.

```
17%6
```

```
mod( 17, 6 )
```

The following value is returned.

netWorkDays

netWorkDays returns the number of working days between the provided two date fields or even custom date values generated using other Quick date functions such as `parseDate` or `epochDate` as an integer.

netWorkDays assumes a standard 5-day work week starting from Monday and ending on Friday. Saturday and Sunday are assumed to be weekends. The calculation is inclusive of both `startDate` and `endDate`. The function operates on and shows results for DAY granularity.

Syntax

```
netWorkDays(startDate, endDate)
```

Arguments

startDate

A valid non-NULL date that acts as the start date for the calculation.

- Dataset fields: Any date field from the dataset that you are adding this function to.
- Date Functions: Any date output from another date function, for example, `parseDate`.
- Calculated fields: Any Quick calculated field that returns a date value.
- Parameters: Any Quick DateTime parameter.
- Any combination of the above stated argument values.

endDate

A valid non-NULL date that acts as the end date for the calculation.

- Dataset fields: Any date field from the dataset that you are adding this function to.
- Date Functions: Any date output from another date function, for example, `parseDate`.
- Calculated fields: Any Quick calculated field that returns a date value.
- Parameters: Any Quick DateTime parameter.
- Any combination of the above stated argument values.

Return type

Integer

Output values

Expected output values include:

- Positive integer (when `start_date < end_date`)
- Negative integer (when `start_date > end_date`)
- NULL when one or both of the arguments get a null value from the `dataset` field.

Example

The following example returns the number of work days falling between two dates.

Let's assume that there's a field named `application_date` with the following values:

```
netWorkDays({startDate}, {endDate})
```

The following are the given field values.

```
startDate endDate netWorkDays
9/4/2022 9/11/2022 5
9/9/2022 9/2/2022 -6
9/10/2022 9/11/2022 0
9/12/2022 9/12/2022 1
```

The following example calculates the number of days worked by each employee and the salary expended per day for each employee:

```
days_worked = netWorkDays({employment_start_date}, {employment_end_date})
salary_per_day = {salary}/{days_worked}
```

The following example filters employees whose employment ends on a work day and determines whether their employment began on work day or a weekend using conditional formatting:

```
is_start_date_work_day = netWorkDays(employment_start_date)
```



```
is_end_date_work_day = netWorkDays(employment_end_date)
```

Field wells

Group by

employee_id

employment_start_date

employment_end_date

salary_paid

Value

days_worked (Sum)

salary_per_day (Sum)

netWorkDays ▼ +

Salary paid to each employee per day

employee_id	employment_start_date	employment_end_date	salary_paid	days_worked	salary_per_day
101	Jan 1, 2021	May 6, 2022	120,000	351	341.88
102	Jan 4, 2021	May 15, 2022	110,000	355	309.86
103	Jan 9, 2021	Aug 3, 2022	105,000	408	257.35
104	Jan 10, 2021	Jan 8, 2022	100,000	260	384.62
105	Aug 3, 2021	Jan 14, 2022	135,000	119	1,134.45
106	Aug 7, 2021	Sep 2, 2022	118,000	280	421.43
107	Aug 9, 2021	Aug 9, 2022	100,000	262	381.68
				2,035	3,231.27

Now

For database datasets that directly query the database, now returns the current date and time using the settings and format specified by the database server. For SPICE and Salesforce data sets, now returns the UTC date and time, in the format yyyy-MM-ddTkk:mm:ss:SSSZ (for example, 2015-10-15T19:11:51:003Z).

Syntax

```
now()
```

Return type

Date

notIn

notIn evaluates if an expression exists within a literal list. If the list doesn't contain the expression, notIn returns true, and otherwise it returns false. notIn is case sensitive for string type inputs.

notIn accepts two kinds of literal list, one is manually entered list and the other is a [multivalue parameter](#).

Syntax

Using a manually entered list:

```
notIn(expression, [literal-1, ...])
```

Using a multivalue parameter:

```
notIn(expression, $multivalue_parameter)
```

Arguments

expression

The expression to be compared with the elements in literal list. It can be a field name like address, a literal value like ' **Unknown**', a single value parameter, or a call to another scalar function—provided this function is not an aggregate function or a table calculation.

literal list

(required) This can be a manually entered list or a multivalue parameter. This argument accepts up to 5,000 elements. However, in a direct query to a third party data source, for example Oracle or Teradata, the restriction can be smaller.

- **manually entered list** – One or more literal values in a list to be compared with the expression. The list should be enclosed in square brackets. All the literals to compare must have the same datatype as the expression.
- **multivalue parameter** – A pre-defined multivalue parameter passed in as a literal list. The multivalue parameter must have the same datatype as the expression.

Return type

Boolean: TRUE/FALSE

Example with a manually entered list

The following example evaluates the `origin_state_name` field for values in a list of string. When comparing string type input, `notIn` only supports case sensitive comparison.

```
notIn(origin_state_name, ["Georgia", "Ohio", "Texas"])
```

The following are the given field values.

```
"Washington"  
  "ohio"  
  "Texas"
```

For these field values the following values are returned.

```
true  
  true  
  false
```

The third return value is false because only "Texas" is one of the excluded values.

The following example evaluates the `f1_date` field for values in a list of string. In order to match the type, `toString` is used to cast the date type to string type.

```
notIn(toString(f1_date), ["2015-05-14", "2015-05-15", "2015-05-16"])
```

Field wells

Group by

fl_date

flightNotInMidMay

Value

Add measures here

Sheet 1 +

FL_date and Flightnotinmidmay

fl_date	flightNotInMidMay
May 15, 2015	0
May 12, 2015	1
May 13, 2015	1

Literals and NULL values are supported in expression argument to be compared with the literals in list. Both of the following two examples will generate a new column of FALSE values.

```
notIn("Washington",["Washington","Ohio"])
```

```
notIn(NULL,[NULL,"Ohio"])
```

Example with multivalue parameter

Let's say an author creates a [multivalue parameter](#) that contains a list of all the state names. Then the author adds a control to allow the reader to select values from the list.

Next, the reader selects three values—"Georgia", "Ohio", and "Texas"—from the parameter's drop down list control. In this case, the following expression is equivalent to the first example, where those three state names are passed as the literal list to be compared with the `original_state_name` field.

```
notIn (origin_state_name, ${stateName MultivalueParameter})
```

Example with ifelse

notIn can be nested in other functions as a boolean value. One example is that authors can evaluate any expression in a list and return the value they want by using notIn and ifelse. The following example evaluates if the dest_state_name of a flight is in a particular list of US states and returns different categories of the states based on the comparison.

```
ifelse(notIn(dest_state_name,["Washington", "Oregon", "California"]),  
      "notWestCoastUSState", "WestCoastUSState")
```

Field wells

Group by

dest_state_nm

StateCategory

Value

Add measures here

Sheet 1

Dest_state_nm and Statecategory

dest_state_nm	StateCategory
Alabama	notWestCoastUSState
Alaska	notWestCoastUSState
California	WestCoastUSState
Colorado	notWestCoastUSState
Connecticut	notWestCoastUSState
Florida	notWestCoastUSState
Georgia	notWestCoastUSState
Hawaii	notWestCoastUSState

nullIf

nullIf compares two expressions. If they are equal, the function returns null. If they are not equal, the function returns the first expression.

Syntax

```
nullIf(expression1, expression2)
```

Arguments

nullIf takes two expressions as arguments.

expression

The expression can be numeric, datetime, or string. It can be a field name, a literal value, or another function.

Return type

String

Example

The following example returns nulls if the reason for a shipment delay is unknown.

```
nullIf(delayReason, 'unknown')
```

The following are the given field values.

```
delayReason
=====
unknown
back ordered
weather delay
```

For these field values, the following values are returned.

```
(null)
back ordered
weather delay
```

parseDate

`parseDate` parses a string to determine if it contains a date value, and returns a standard date in the format `yyyy-MM-ddTkk:mm:ss.SSSZ` (using the format pattern syntax specified in [Class DateTimeFormat](#) in the Joda project documentation), for example `2015-10-15T19:11:51.003Z`. This function returns all rows that contain a date in a valid format and skips any rows that don't, including rows that contain null values.

Quick supports dates in the range from Jan 1, 1900 00:00:00 UTC to Dec 31, 2037 23:59:59 UTC. For more information, see [Supported date formats](#).

Syntax

```
parseDate(expression, ['format'])
```

Arguments

expression

The expression must be a string. It can be the name of a field that uses the string data type, a literal value like `'1/1/2016'`, or a call to another function that outputs a string.

format

(Optional) A string containing the format pattern that *date_string* must match. For example, if you are using a field with data like `01/03/2016`, you specify the format `'MM/dd/yyyy'`. If you don't specify a format, it defaults to `yyyy-MM-dd`. Rows whose data doesn't conform to *format* are skipped.

Different date formats are supported based on the type of dataset used. Use the following table to see details of supported date formats.

Date source type	Supported date formats
File, Amazon Athena, and Salesforce datasets	All date format patterns specified in Supported date formats .
Direct query of Amazon Aurora, MariaDB, and MySQL databases	- MM/dd/yyyy

Date source type	Supported date formats
	dd/MM/yyyy • yyyy/MM/dd • MMM/dd/yyyy • dd/MMM/yyyy • yyyy/MMM/dd • MM/dd/yyyy HH:mm:ss • dd/MM/yyyy HH:mm:ss • yyyy/MM/dd HH:mm:ss • MMM/dd/yyyy HH:mm:ss • dd/MMM/yyyy HH:mm:ss • yyyy/MMM/dd HH:mm:ss • MM-dd-yyyy • dd-MM-yyyy • yyyy-MM-dd • MMM-dd-yyyy • dd-MMM-yyyy • yyyy-MMM-dd • MM-dd-yyyy HH:mm:ss •

Date source type	Supported date formats
	dd-MM-yyyy HH:mm:ss • yyyy-MM-dd HH:mm:ss • MMM-dd-yyyy HH:mm:ss • dd-MMM-yyyy HH:mm:ss • yyyy-MMM-dd HH:mm:ss • MM/dd/yyyy HH:mm:ss.SSS • dd/MM/yyyy HH:mm:ss.SSS • yyyy/MM/dd HH:mm:ss.SSS • MMM/dd/yyyy HH:mm:ss.SSS • dd/MMM/yyyy HH:mm:ss.SSS • yyyy/MMM/dd HH:mm:ss.SSS • MM-dd-yyyy HH:mm:ss.SSS • dd-MM-yyyy HH:mm:ss.SSS • yyyy-MM-dd HH:mm:ss.SSS • MMM-dd-yyyy HH:mm:ss.SSS • dd-MMM-yyyy HH:mm:ss.SSS • yyyy-MMM-dd HH:mm:ss.SSS

Date source type	Supported date formats
Direct query of Snowflake	• dd/MM/yyyy • dd/MM/yyyy HH:mm:ss • dd-MM-yyyy • dd-MM-yyyy HH:mm:ss • MM/dd/yyyy • MM/dd/yyyy HH:mm:ss • MM-dd-yyyy • MM-dd-yyyy HH:mm:ss • yyyy/MM/dd • yyyy/MM/dd HH:mm:ss • yyyy-MM-dd • yyyy-MM-dd HH:mm:ss • MM/dd/yyyy HH:mm:ss.SSS • dd/MM/yyyy HH:mm:ss.SSS • yyyy/MM/dd HH:mm:ss.SSS • MMM/dd/yyyy HH:mm:ss.SSS • dd/MMM/yyyy HH:mm:ss.SSS • yyyy/MMM/dd HH:mm:ss.SSS •

Date source type	Supported date formats
	MM-dd-yyyy HH:mm:ss.SSS • dd-MM-yyyy HH:mm:ss.SSS • yyyy-MM-dd HH:mm:ss.SSS • MMM-dd-yyyy HH:mm:ss.SSS • dd-MMM-yyyy HH:mm:ss.SSS • yyyy-MMM-dd HH:mm:ss.SSS

Date source type	Supported date formats
Direct query of Microsoft SQL Server databases	• dd-MM-yyyy • MM/dd/yyyy • dd/MM/yyyy • yyyy/MM/dd • MMM/dd/yyyy • dd/MMM/yyyy • yyyy/MMM/dd • dd/MM/yyyy HH:mm:ss • yyyy/MM/dd HH:mm:ss • MMM/dd/yyyy HH:mm:ss • dd/MMM/yyyy HH:mm:ss • yyyy/MMM/dd HH:mm:ss • MM-dd-yyyy • yyyy-MM-dd • MMM-dd-yyyy • yyyy-MMM-dd • MM-dd-yyyy HH:mm:ss • dd-MM-yyyy HH:mm:ss •

Date source type	Supported date formats
	yyyy-MM-dd HH:mm:ss • MMM-dd-yyyy HH:mm:ss • dd-MMM-yyyy HH:mm:ss • yyyy-MMM-dd HH:mm:ss • MM/dd/yyyy HH:mm:ss.SSS • dd/MM/yyyy HH:mm:ss.SSS • yyyy/MM/dd HH:mm:ss.SSS • MMM/dd/yyyy HH:mm:ss.SSS • dd/MMM/yyyy HH:mm:ss.SSS • yyyy/MMM/dd HH:mm:ss.SSS • MM-dd-yyyy HH:mm:ss.SSS • dd-MM-yyyy HH:mm:ss.SSS • yyyy-MM-dd HH:mm:ss.SSS • MMM-dd-yyyy HH:mm:ss.SSS • dd-MMM-yyyy HH:mm:ss.SSS • yyyy-MMM-dd HH:mm:ss.SSS

Date source type	Supported date formats
Direct query of Amazon Redshift or PostgreSQL databases Also, datasets from any DBMS that are stored in Quick SPICE	• MM/dd/yyyy • dd/MM/yyyy • yyyy/MM/dd • MMM/dd/yyyy • dd/MMM/yyyy • yyyy/MMM/dd • MM/dd/yyyy HH:mm:ss • dd/MM/yyyy HH:mm:ss • yyyy/MM/dd HH:mm:ss • MMM/dd/yyyy HH:mm:ss • dd/MMM/yyyy HH:mm:ss • yyyy/MMM/dd HH:mm:ss • MM-dd-yyyy • dd-MM-yyyy • yyyy-MM-dd • MMM-dd-yyyy • dd-MMM-yyyy • yyyy-MMM-dd •

Date source type	Supported date formats
	MM-dd-yyyy HH:mm:ss • dd-MM-yyyy HH:mm:ss • yyyy-MM-dd HH:mm:ss • MMM-dd-yyyy HH:mm:ss • dd-MMM-yyyy HH:mm:ss • yyyy-MMM-dd HH:mm:ss • yyyyMMdd'T'HHmmss • yyyy-MM-dd'T'HH:mm:ss • MM/dd/yyyy HH:mm:ss.SSS • dd/MM/yyyy HH:mm:ss.SSS • yyyy/MM/dd HH:mm:ss.SSS • MMM/dd/yyyy HH:mm:ss.SSS • dd/MMM/yyyy HH:mm:ss.SSS • yyyy/MMM/dd HH:mm:ss.SSS • MM-dd-yyyy HH:mm:ss.SSS • dd-MM-yyyy HH:mm:ss.SSS • yyyy-MM-dd HH:mm:ss.SSS • MMM-dd-yyyy HH:mm:ss.SSS •

Date source type	Supported date formats
	dd-MMM-yyyy HH:mm:ss.SSS • yyyy-MMM-dd HH:mm:ss.SSS

Return type

Date

Example

The following example evaluates `prodDate` to determine if it contains date values.

```
parseDate(prodDate, 'MM/dd/yyyy')
```

The following are the given field values.

```
prodDate
-----
01-01-1999
12/31/2006
1/18/1982
7/4/2010
```

For these field values, the following rows are returned.

```
12-31-2006T00:00:00.000Z
01-18-1982T00:00:00.000Z
07-04-2010T00:00:00.000Z
```

parseDecimal

`parseDecimal` parses a string to determine if it contains a decimal value. This function returns all rows that contain a decimal, integer, or null value, and skips any rows that don't. If the row contains an integer value, it is returned as a decimal with up to 4 decimal places. For example, a value of '2' is returned as '2.0'.

Syntax

```
parseDecimal(expression)
```

Arguments

expression

The expression must be a string. It can be the name of a field that uses the string data type, a literal value like **'9.62'**, or a call to another function that outputs a string.

Return type

Decimal(Fixed) in the legacy data preparation experience.

Decimal(Float) in the new data preparation experience.

Example

The following example evaluates fee to determine if it contains decimal values.

```
parseDecimal(fee)
```

The following are the given field values.

```
fee
-----
2
2a
12.13
3b
3.9
(null)
198.353398
```

For these field values, the following rows are returned.

```
2.0
12.13
3.9
```

```
(null)
198.3533
```

parseInt

`parseInt` parses a string to determine if it contains an integer value. This function returns all rows that contain a decimal, integer, or null value, and skips any rows that don't. If the row contains a decimal value, it is returned as the nearest integer, rounded down. For example, a value of '2.99' is returned as '2'.

Syntax

```
parseInt(expression)
```

Arguments

expression

The expression must be a string. It can be the name of a field that uses the string data type, a literal value like '3', or a call to another function that outputs a string.

Return type

Integer

Example

The following example evaluates `feeType` to determine if it contains integer values.

```
parseInt(feeType)
```

The following are the given field values.

```
feeType
-----
2
2.1
2a
3
```

```
3b
(null)
5
```

For these field values, the following rows are returned.

```
2
2
3
(null)
5
```

parseJson

Use `parseJson` to extract values from a JSON object.

If your dataset is stored in Quick SPICE, you can use `parseJson` when you are preparing a data set, but not in calculated fields during analysis.

For direct query, you can use `parseJson` both during data preparation and analysis. The `parseJson` function applies to either strings or to JSON native data types, depending on the dialect, as shown in the following table.

Dialect	Type
PostgreSQL	JSON
Amazon Redshift	String
Microsoft SQL Server	String
MySQL	JSON
Teradata	JSON
Oracle	String
Presto	String
Snowflake	Semistructured data type object and array

Dialect	Type
Hive	String

Syntax

```
parseJson(fieldName, path)
```

Arguments

fieldName

The field containing the JSON object that you want to parse.

path

The path to the data element you want to parse from the JSON object. Only letters, numbers, and blank spaces are supported in the path argument. Valid path syntax includes:

- `$` – Root object
- `.` – Child operator
- `[]` – Subscript operator for array

Return type

String

Example

The following example evaluates incoming JSON to retrieve a value for item quantity. By using this during data preparation, you can create a table out of the JSON.

```
parseJson({jsonField}, "$.items.qty")
```

The following shows the JSON.

```
{  
  "customer": "John Doe",
```

```
"items": {
  "product": "Beer",
  "qty": 6
},
"list1": [
  "val1",
  "val2"
],
"list2": [
  {
    "list21key1": "list1value1"
  }
]
}
```

For this example, the following value is returned.

```
6
```

Example

The following example evaluates `JSONObject1` to extract the first key value pair (KVP), labeled "State", and assign the value to the calculated field that you are creating.

```
parseJson(JSONObject1, "$.state")
```

The following are the given field values.

```
JSONObject1
-----
{"State":"New York","Product":"Produce","Date Sold":"1/16/2018","Sales Amount":"$3423.39"}
{"State":"North Carolina","Product":"Bakery Products","Date Sold":"2/1/2018","Sales Amount":"$3226.42"}
{"State":"Utah","Product":"Water","Date Sold":"4/24/2018","Sales Amount":"$7001.52"}
```

For these field values, the following rows are returned.

```
New York
North Carolina
```

```
Utah
```

Replace

`replace` replaces part of a string with another string that you specify.

Syntax

```
replace(expression, substring, replacement)
```

Arguments

expression

The expression must be a string. It can be the name of a field that uses the string data type, a literal value like **'12 Main Street'**, or a call to another function that outputs a string.

substring

The set of characters in *expression* that you want to replace. The substring can occur one or more times in *expression*.

replacement

The string you want to have substituted for *substring*.

Return type

String

Example

The following example replaces the substring 'and' with 'or'.

```
replace('1 and 2 and 3', 'and', 'or')
```

The following string is returned.

```
1 or 2 or 3
```

Right

`right` returns the rightmost characters from a string, including spaces. You specify the number of characters to be returned.

Syntax

```
right(expression, limit)
```

Arguments

expression

The expression must be a string. It can be the name of a field that uses the string data type, a literal value like **'12 Main Street'**, or a call to another function that outputs a string.

limit

The number of characters to be returned from *expression*, starting from the last character in the string.

Return type

String

Example

The following example returns the last five characters from a string.

```
right('Seattle Store#14', 12)
```

The following value is returned.

```
tle Store#14
```

Round

`round` rounds a decimal value to the closest integer if no scale is specified, or to the closest decimal place if scale is specified.

Syntax

```
round(decimal, scale)
```

Arguments

decimal

A field that uses the decimal data type, a literal value like **17.62**, or a call to another function that outputs a decimal.

scale

The number of decimal places to use for the return values.

Return type

Operand	Return type in the legacy data preparation experience	Return type in the new data preparation experience
INT	DECIMAL(FIXED)	DECIMAL(FIXED)
DECIMAL(FIXED)	DECIMAL(FIXED)	DECIMAL(FIXED)
DECIMAL(FLOAT)	DECIMAL(FIXED)	DECIMAL(FLOAT)

Example

The following example rounds a decimal field to the closest second decimal place.

```
round(salesAmount, 2)
```

The following are the given field values.

```
20.1307  
892.0388
```



```
57.5447
```

For these field values, the following values are returned.

```
20.13  
892.04  
58.54
```

Rtrim

`rtrim` removes following blank space from a string.

Syntax

```
rtrim(expression)
```

Arguments

expression

The expression must be a string. It can be the name of a field that uses the string data type, a literal value like **'12 Main Street'**, or a call to another function that outputs a string.

Return type

String

Example

The following example removes the following spaces from a string.

```
rtrim('Seattle Store #14  ')
```

For these field values, the following values are returned.

```
Seattle Store #14
```

Split

`split` splits a string into an array of substrings, based on a delimiter that you choose, and returns the item specified by the position.

You can only add `split` to a calculated field during data preparation, not to an analysis. This function is not supported in direct queries to Microsoft SQL Server.

Syntax

```
split(expression, delimiter , position)
```

Arguments

expression

The expression must be a string. It can be the name of a field that uses the string data type, a literal value like `'12 Main Street;1402 35th Ave;1818 Elm Ct;11 Janes Lane'`, or a call to another function that outputs a string.

delimiter

The character that delimits where the string is broken into substrings. For example, `split('one|two|three', '|', 2)` becomes the following.

```
one  
two  
three
```

If you choose `position = 2`, `split` returns `'two'`.

position

(Required) The position of the item to return from the array. The position of the first item in the array is 1.

Return type

String array

Example

The following example splits a string into an array, using the semicolon character (;) as the delimiter, and returns the third element of the array.

```
split('123 Test St;1402 35th Ave;1818 Elm Ct;11 Janes Lane', ';', 3)
```

The following item is returned.

```
1818 Elm Ct
```

This function skips items containing null values or empty strings.

Sqrt

`sqrt` returns the square root of a given expression.

Syntax

```
sqrt(expression)
```

Arguments

expression

The expression must be numeric. It can be a field name, a literal value, or another function.

startsWith

`startsWith` evaluates if the expression starts with a substring that you specify. If the expression starts with the substring, `startsWith` returns true, and otherwise it returns false.

Syntax

```
startsWith(expression, substring, string-comparison-mode)
```

Arguments

expression

The expression must be a string. It can be the name of a field that uses the string data type, a literal value like **'12 Main Street'**, or a call to another function that outputs a string.

substring

The set of characters to check against the *expression*. The substring can occur one or more times in the *expression*.

string-comparison-mode

(Optional) Specifies the string comparison mode to use:

- `CASE_SENSITIVE` – String comparisons are case-sensitive.
- `CASE_INSENSITIVE` – String comparisons are case-insensitive.

This value defaults to `CASE_SENSITIVE` when blank.

Return type

Boolean

Examples

Default case sensitive example

The following case sensitive example evaluates if `state_nm` startsWith **New**.

```
startsWith(state_nm, "New")
```

The following are the given field values.

```
New York  
new york
```

For these field values, the following values are returned.

```
true  
false
```

Case insensitive example

The following case insensitive example evaluates if `state_nm` startsWith **new**.

```
startsWith(state_nm, "new", CASE_INSENSITIVE)
```

The following are the given field values.

```
New York
```

```
new york
```

For these field values, the following values are returned.

```
true  
true
```

Example with conditional statements

The `startsWith` function can be used as the conditional statement within the following If functions: [avgIf](#), [minIf](#), [distinct_countIf](#), [countIf](#), [maxIf](#), [medianIf](#), [stdevIf](#), [stdevplf](#), [sumIf](#), [varIf](#), and [varplf](#).

The following example sums Sales only if `state_nm` starts with **New**.

```
sumIf(Sales, startsWith(state_nm, "New"))
```

Does NOT contain example

The conditional NOT operator can be used to evaluate if the expression does not start with the specified substring.

```
NOT(startsWith(state_nm, "New"))
```

Example using numeric values

Numeric values can be used in the expression or substring arguments by applying the `toString` function.

```
startsWith(state_nm, toString(5) )
```

Strlen

`strlen` returns the number of characters in a string, including spaces.

Syntax

```
strlen(expression)
```

Arguments

expression

An expression can be the name of a field that uses the string data type like **address1**, a literal value like **'Unknown'**, or another function like `substring(field_name,0,5)`.

Return type

Integer

Example

The following example returns the length of the specified string.

```
strlen('1421 Main Street')
```

The following value is returned.

```
16
```

Substring

`substring` returns the characters in a string, starting at the location specified by the *start* argument and proceeding for the number of characters specified by the *length* arguments.

Syntax

```
substring(expression, start, length)
```

Arguments

expression

An expression can be the name of a field that uses the string data type like **address1**, a literal value like **'Unknown'**, or another function like `substring(field_name,1,5)`.

start

The character location to start from. *start* is inclusive, so the character at the starting position is the first character in the returned value. The minimum value for *start* is 1.

length

The number of additional characters to include after *start*. *length* is inclusive of *start*, so the last character returned is (*length* - 1) after the starting character.

Return type

String

Example

The following example returns the 13th through 19th characters in a string. The beginning of the string is index 1, so you begin counting at the first character.

```
substring('Fantasy and Science Fiction',13,7)
```

The following value is returned.

```
Science
```

switch

`switch` compares a *condition-expression* with the literal labels, within a set of literal label and *return-expression* pairings. It then returns the *return-expression* corresponding to the first literal label that's equal to the *condition-expression*. If no label equals to the *condition-expression*, `switch` returns the *default-expression*. Every *return-expression* and *default-expression* must have the same datatype.

Syntax

```
switch(condition-expression, label-1, return-expression-1 [, label-n, return-expression-n ...],  
      default-expression)
```

Arguments

`switch` requires one or more *if,then* expression pairings, and requires exactly one expression for the *else* argument.

condition-expression

The expression to be compared with the label-literals. It can be a field name like `address`, a literal value like `'Unknown'`, or another scalar function like `toString(salesAmount)`.

label

The literal to be compared with the *condition-expression* argument, all of the literals must have the same data type as *condition-expression* argument. `switch` accepts up to 5000 labels.

return-expression

The expression to return if the value of its label equals to the value of the *condition-expression*. It can be a field name like `address`, a literal value like `'Unknown'`, or another scalar function like `toString(salesAmount)`. All of the *return-expression* arguments must have the same data type as the *default-expression*.

default-expression

The expression to return if no value of any label arguments equals to the value of *condition-expression*. It can be a field name like `address`, a literal value like `'Unknown'`, or another scalar function like `toString(salesAmount)`. The *default-expression* must have the same data type as all of the *return-expression* arguments.

Return type

`switch` returns a value of the same data type as the values in *return-expression*. All data returned *return-expression* and *default-expression* must be of the same data type or be converted to the same data type.

General Examples

The following example returns the AWS Region code of input region name.

```
switch(region_name,  
      "US East (N. Virginia)", "us-east-1",  
      "Europe (Ireland)", "eu-west-1",  
      "US West (N. California)", "us-west-1",  
      "other regions")
```

The following are the given field values.


```
"US East (N. Virginia)"
  "US West (N. California)"
  "Asia Pacific (Tokyo)"
```

For these field values the following values are returned.

```
"us-east-1"
  "us-west-1"
  "other regions"
```

Use switch to replace ifelse

The following `ifelse` use case is an equivalent of the previous example, for `ifelse` evaluating whether values of one field equals to different literal values, using `switch` instead is a better choice.

```
ifelse(region_name = "US East (N. Virginia)", "us-east-1",
       region_name = "Europe (Ireland)", "eu-west-1",
       region_name = "US West (N. California)", "us-west-1",
       "other regions")
```

Expression as return value

The following example uses expressions in *return-expressions*:

```
switch({origin_city_name},
      "Albany, NY", {arr_delay} + 20,
      "Alexandria, LA", {arr_delay} - 10,
      "New York, NY", {arr_delay} * 2,
      {arr_delay})
```

The preceding example changes the expected delay time for each flight from a particular city.

Field wells

Group by

origin_city_name



arr_delay



newDelayToSomeCities



Value

Add measures here

Sheet 1



Origin_city_name, Arr_delay, and Newdelaytosomecities

origin_city_name	arr_delay	newDelayToSomeCities
New York, NY	-33	-66
New York, NY	-23	-46
New York, NY	-20	-40
New York, NY	-14	-28
Alexandria, LA	-11	-21
New York, NY	-9	-18
New York, NY	-1	-2
Albany, NY	-20	0
Albany, NY	-5	15
New York, NY	9	18
Albany, NY	14	34
New York, NY	20	40

toLower

toLower formats a string in all lowercase. toLower skips rows containing null values.

Syntax

```
toLower(expression)
```

Arguments

expression

The expression must be a string. It can be the name of a field that uses the string data type, a literal value like **'12 Main Street'**, or a call to another function that outputs a string.

Return type

String

Example

The following example converts a string value into lowercase.

```
toLower('Seattle Store #14')
```

The following value is returned.

```
seattle store #14
```

toString

toString formats the input expression as a string. toString skips rows containing null values.

Syntax

```
toString(expression)
```

Arguments

expression

An expression can be a field of any data type, a literal value like **14.62**, or a call to another function that returns any data type.

Return type

String

Example

The following example returns the values from `payDate` (which uses the date data type) as strings.

```
toString(payDate)
```

The following are the given field values.

```
payDate
-----
1992-11-14T00:00:00.000Z
2012-10-12T00:00:00.000Z
1973-04-08T00:00:00.000Z
```

For these field values, the following rows are returned.

```
1992-11-14T00:00:00.000Z
2012-10-12T00:00:00.000Z
1973-04-08T00:00:00.000Z
```

toUpper

`toUpper` formats a string in all uppercase. `toUpper` skips rows containing null values.

Syntax

```
toUpper(expression)
```

Arguments

expression

The expression must be a string. It can be the name of a field that uses the string data type, a literal value like **'12 Main Street'**, or a call to another function that outputs a string.

Return type

String

Example

The following example converts a string value into uppercase.

```
toUpper('Seattle Store #14')
```

The following value is returned.

```
SEATTLE STORE #14
```

trim

`trim` removes both preceding and following blank space from a string.

Syntax

```
trim(expression)
```

Arguments

expression

The expression must be a string. It can be the name of a field that uses the string data type, a literal value like **'12 Main Street'**, or a call to another function that outputs a string.

Return type

String

Example

The following example removes the following spaces from a string.

```
trim('  Seattle Store #14  ')
```

For these field values, the following values are returned.

```
Seattle Store #14
```

truncDate

`truncDate` returns a date value that represents a specified portion of a date. For example, requesting the year portion of the value `2012-09-02T00:00:00.000Z` returns `2012-01-01T00:00:00.000Z`. Specifying a time-related period for a date that doesn't contain time information returns the initial date value unchanged.

Syntax

```
truncDate('period', date)
```

Arguments

period

The period of the date that you want returned. Valid periods are as follows:

- `YYYY`: This returns the year portion of the date.
- `Q`: This returns the date of the first day of the quarter that the date belongs to.
- `MM`: This returns the month portion of the date.
- `DD`: This returns the day portion of the date.
- `WK`: This returns the week portion of the date. The week starts on Sunday in Amazon Quick.
- `HH`: This returns the hour portion of the date.
- `MI`: This returns the minute portion of the date.
- `SS`: This returns the second portion of the date.
- `MS`: This returns the millisecond portion of the date.

date

A date field or a call to another function that outputs a date.

Return type

Date

Example

The following example returns a date representing the month of the order date.

```
truncDate('MM', orderDate)
```

The following are the given field values.

```
orderDate
=====
2012-12-14T00:00:00.000Z
2013-12-29T00:00:00.000Z
2012-11-15T00:00:00.000Z
```

For these field values, the following values are returned.

```
2012-12-01T00:00:00.000Z
2013-12-01T00:00:00.000Z
2012-11-01T00:00:00.000Z
```

Aggregate functions

Aggregate functions are only available during analysis and visualization. Each of these functions returns values grouped by the chosen dimension or dimensions. For each aggregation, there is also a conditional aggregation. These perform the same type of aggregation, based on a condition.

When a calculated field formula contains an aggregation, it becomes a custom aggregation. To make sure that your data is accurately displayed, Amazon Quick applies the following rules:

- Custom aggregations can't contain nested aggregate functions. For example, this formula doesn't work: `sum( avg(x)/avg(y) )`. However, nesting nonaggregated functions inside or outside aggregate functions does work. For example, `ceil( avg(x) )` works. So does `avg( ceil(x) )`.
- Custom aggregations can't contain both aggregated and nonaggregated fields, in any combination. For example, this formula doesn't work: `Sum(sales)+quantity`.
- Filter groups can't contain both aggregated and nonaggregated fields.
- Custom aggregations can't be converted to a dimension. They also can't be dropped into the field well as a dimension.
- In a pivot table, custom aggregations can't be added to table calculations.
- Scatter plots with custom aggregations need at least one dimension under **Group/Color** in the field wells.

For more information about supported functions and operators, see [Calculated field function and operator reference for Amazon Quick](#).

The aggregate functions for calculated fields in Quick include the following.

Topics

- [avg](#)
- [avgIf](#)
- [count](#)
- [countIf](#)
- [distinct_count](#)
- [distinct_countIf](#)
- [max](#)
- [maxIf](#)
- [median](#)
- [medianIf](#)
- [min](#)
- [minIf](#)
- [percentile](#)
- [percentileCont](#)
- [percentileDisc \(percentile\)](#)
- [periodToDateAvg](#)
- [periodToDateCount](#)
- [periodToDateMax](#)
- [periodToDateMedian](#)
- [periodToDateMin](#)
- [periodToDatePercentile](#)
- [periodToDatePercentileCont](#)
- [periodToDateStDev](#)
- [periodToDateStDevP](#)
- [periodToDateSum](#)
- [periodToDateVar](#)
- [periodToDateVarP](#)
- [stdev](#)

- [stdevp](#)
- [stdevlf](#)
- [stdevplf](#)
- [sum](#)
- [sumlf](#)
- [var](#)
- [varlf](#)
- [varp](#)
- [varplf](#)

avg

The avg function averages the set of numbers in the specified measure, grouped by the chosen dimension or dimensions. For example, `avg(salesAmount)` returns the average for that measure grouped by the (optional) chosen dimension.

Syntax

```
avg(decimal, [group-by level])
```

Arguments

decimal

The argument must be a measure. Null values are omitted from the results. Literal values don't work. The argument must be a field.

group-by level

(Optional) Specifies the level to group the aggregation by. The level added can be any dimension or dimensions independent of the dimensions added to the visual.

The argument must be a dimension field. The group-by level must be enclosed in square brackets []. For more information, see [Level-aware calculation - aggregate \(LAC-A\) functions](#).

Examples

The following example calculates the average sales.

```
avg({Sales})
```

You can also specify at what level to group the computation using one or more dimensions in the view or in your dataset. This is called a LAC-A function. For more information about LAC-A functions, see [Level-aware calculation - aggregate \(LAC-A\) functions](#). The following example calculates the average sales at the Country level, but not across other dimensions (Region or Product) in the visual.

```
avg({Sales}, [{Country}])
```

Avg(sales), Sum of Avg(sales and [Country]) by Region, Country, and Product

Regi...	Country	Product	avg(sales)	avg(sales, [Country])
AMER	Argentina	Big Ol Database	3,299.95	275.11
AMER	Argentina	ChatBot Plugin	67.53	275.11
AMER	Argentina	ContactMatcher	232.22	275.11
AMER	Argentina	Data Smasher	113.73	275.11
AMER	Argentina	FinanceHub	209.86	275.11
AMER	Argentina	Marketing Suite	1,137.94	275.11
AMER	Argentina	Marketing Suite - Gold	518.79	275.11
AMER	Argentina	OneView	323.41	275.11
AMER	Argentina	SaaS Connector Pack	67.93	275.11
AMER	Argentina	SaaS Connector Pack - Gold	38.43	275.11
AMER	Argentina	Site Analytics	397.53	275.11
AMER	Argentina	Storage	10.82	275.11
AMER	Argentina	Support	73.15	275.11
APJ	Australia	Alchemy	1,479.97	162.94
APJ	Australia	Big Ol Database	751.26	162.94
APJ	Australia	ChatBot Plugin	25.14	162.94
APJ	Australia	ContactMatcher	113.16	162.94
API	Australia	Data Smasher	167.75	162.94

avgIf

Based on a conditional statement, the avgIf function averages the set of numbers in the specified measure, grouped by the chosen dimension or dimensions. For example, `avgIf(ProdRev, CalendarDay >= ${BasePeriodStartDate} AND CalendarDay <= ${BasePeriodEndDate} AND SourcingType <> 'Indirect')` returns the average for that measure grouped by the (optional) chosen dimension, if the condition evaluates to true.

Syntax

```
avgIf(dimension or measure, condition)
```

Arguments

decimal

The argument must be a measure. Null values are omitted from the results. Literal values don't work. The argument must be a field.

condition

One or more conditions in a single statement.

count

The count function calculates the number of values in a dimension or measure, grouped by the chosen dimension or dimensions. For example, `count(product type)` returns the total number of product types grouped by the (optional) chosen dimension, including any duplicates. The `count(sales)` function returns the total number of sales completed grouped by the (optional) chosen dimension, for example salesperson.

Syntax

```
count(dimension or measure, [group-by level])
```

Arguments

dimension or measure

The argument must be a measure or a dimension. Null values are omitted from the results. Literal values don't work. The argument must be a field.

group-by level


(Optional) Specifies the level to group the aggregation by. The level added can be any dimension or dimensions independent of the dimensions added to the visual.

The argument must be a dimension field. The group-by level must be enclosed in square brackets []. For more information, see [Level-aware calculation - aggregate \(LAC-A\) functions](#).

Examples

The following example calculates the count of sales by a specified dimension in the visual. In this example, the count of sales by month are shown.

```
count({Sales})
```



You can also specify at what level to group the computation using one or more dimensions in the view or in your dataset. This is called a LAC-A function. For more information about LAC-A functions, see [Level-aware calculation - aggregate \(LAC-A\) functions](#). The following example calculates the count of sales at the Country level, but not across other dimensions (Region or Product) in the visual.

```
count({Sales}, [{Country}])
```

Count(sales), Sum of Count(sales and [Country]) by Region, Country, and Product

Regi...	Country	Product	count(Sales)	count(Sales, [Country])
AMER	Argentina	Big OL Database	3	130
AMER	Argentina	ChatBot Plugin	11	130
AMER	Argentina	ContactMatcher	17	130
AMER	Argentina	Data Smasher	9	130
AMER	Argentina	FinanceHub	13	130
AMER	Argentina	Marketing Suite	2	130
AMER	Argentina	Marketing Suite - Gold	9	130
AMER	Argentina	OneView	13	130
AMER	Argentina	SaaS Connector Pack	14	130
AMER	Argentina	SaaS Connector Pack - Gold	4	130
AMER	Argentina	Site Analytics	9	130
AMER	Argentina	Storage	5	130
AMER	Argentina	Support	21	130
APJ	Australia	Alchemy	4	492
APJ	Australia	Big OL Database	5	492
APJ	Australia	ChatBot Plugin	37	492
APJ	Australia	ContactMatcher	98	492

countIf

Based on a conditional statement, the countIf function calculates the number of values in a dimension or measure, grouped by the chosen dimension or dimensions.

Syntax

```
countIf(dimension or measure, condition)
```

Arguments

dimension or measure

The argument must be a measure or a dimension. Null values are omitted from the results. Literal values don't work. The argument must be a field.

condition

One or more conditions in a single statement.

Return type

Integer

Example

The following function returns a count of the sales transactions (Revenue) that meet the conditions, including any duplicates.

```
countIf (  
    Revenue,  
    # Conditions  
        CalendarDay >= ${BasePeriodStartDate} AND  
        CalendarDay <= ${BasePeriodEndDate} AND  
        SourcingType <> 'Indirect'  
)
```

distinct_count

The `distinct_count` function calculates the number of distinct values in a dimension or measure, grouped by the chosen dimension or dimensions. For example, `distinct_count(product type)` returns the total number of unique product types grouped by the (optional) chosen dimension, without any duplicates. The `distinct_count(ship date)` function returns the total number of dates when products were shipped grouped by the (optional) chosen dimension, for example region.

Syntax

```
distinct_count(dimension or measure, [group-by level])
```

Arguments

dimension or measure

The argument must be a measure or a dimension. Null values are omitted from the results. Literal values don't work. The argument must be a field.

group-by level


(Optional) Specifies the level to group the aggregation by. The level added can be any dimension or dimensions independent of the dimensions added to the visual.

The argument must be a dimension field. The group-by level must be enclosed in square brackets []. For more information, see [Level-aware calculation - aggregate \(LAC-A\) functions](#).

Example

The following example calculates the total number of dates when products were ordered grouped by the (optional) chosen dimension in the visual, for example region.

```
distinct_count({Order Date})
```



You can also specify at what level to group the computation using one or more dimensions in the view or in your dataset. This is called a LAC-A function. For more information about LAC-A functions, see [Level-aware calculation - aggregate \(LAC-A\) functions](#). The following example calculates the average sales at the Country level, but not across other dimensions (Region) in the visual.

```
distinct_count({Order Date}, [Country])
```

Distinct_count(order Date), Sum of Distinct_count(order Date and [Country] by Region

Regi...	distinct_count(Order Date)	distinct_count(Order Date, [Country]
AMER	897	1,391
APJ	679	932
EMEA	975	1,885

distinct_countIf

Based on a conditional statement, the `distinct_countIf` function calculates the number of distinct values in a dimension or measure, grouped by the chosen dimension or dimensions. For example, `distinct_countIf(product type)` returns the total number of unique product types grouped by the (optional) chosen dimension, without any duplicates. The `distinct_countIf(ProdRev, CalendarDay >= ${BasePeriodStartDate} AND CalendarDay <= ${BasePeriodEndDate} AND SourcingType <> 'Indirect')` function returns the total number of dates when products were shipped grouped by the (optional) chosen dimension, for example region, if the condition evaluates to true.

Syntax

```
distinct_countIf(dimension or measure, condition)
```

Arguments

dimension or measure

The argument must be a measure or a dimension. Null values are omitted from the results. Literal values don't work. The argument must be a field.

condition

One or more conditions in a single statement.

max

The `max` function returns the maximum value of the specified measure or date, grouped by the chosen dimension or dimensions. For example, `max(sales goal)` returns the maximum sales goals grouped by the (optional) chosen dimension.

Syntax

```
max(measure, [group-by level])
```

Arguments

measure

The argument must be a measure or a date. Null values are omitted from the results. Literal values don't work. The argument must be a field.

Maximum dates work only in the **Value** field well of tables and pivot tables.

group-by level

(Optional) Specifies the level to group the aggregation by. The level added can be any dimension or dimensions independent of the dimensions added to the visual.

The argument must be a dimension field. The group-by level must be enclosed in square brackets []. For more information, see [Level-aware calculation - aggregate \(LAC-A\) functions](#).

Examples

The following example returns the max sales value for each region. It is compared to the total, minimum, and median sales values.

```
max({Sales})
```

Field wells

Group by

Region

Value

Sales (Sum)

max({Sales}) (Custom)

min({Sales}) (Custom)

median({Sales}) (Custom)

Sheet 1

+

Max({sales}), Sum of Sales, Min({sales}), and Median({sales}) by Region

Region	Sales	max({Sales})	min({Sales})	median({Sales})
AMER	837,849.6	13,999.96	0.99	60.34
APJ	415,464.24	17,499.95	0.44	44.43
EMEA	1,043,887.02	22,638.48	0.85	53.42

You can also specify at what level to group the computation using one or more dimensions in the view or in your dataset. This is called a LAC-A function. For more information about LAC-A functions, see [Level-aware calculation - aggregate \(LAC-A\) functions](#). The following example calculates the max sales at the Country level, but not across other dimensions (Region) in the visual.

max({Sales}, [Country])

Max(sales), Sum of Max(sales and [Country]) by Region

Regi...	max(Sales)	max(Sales, [Country])
AMER	13,999.96	51,730.23
APJ	17,499.95	47,417.48
EMEA	22,638.48	128,301.12

maxIf

Based on a conditional statement, the maxIf function returns the maximum value of the specified measure, grouped by the chosen dimension or dimensions. For example, `maxIf(ProdRev, CalendarDay >= ${BasePeriodStartDate} AND CalendarDay <= ${BasePeriodEndDate} AND SourcingType <> 'Indirect')` returns the maximum sales goals grouped by the (optional) chosen dimension, if the condition evaluates to true.

Syntax

```
maxIf(measure, condition)
```

Arguments

measure

The argument must be a measure. Null values are omitted from the results. Literal values don't work. The argument must be a field.

condition

One or more conditions in a single statement.

median

The median aggregation returns the median value of the specified measure, grouped by the chosen dimension or dimensions. For example, `median(revenue)` returns the median revenue grouped by the (optional) chosen dimension.

Syntax

```
median(measure, [group-by level])
```

Arguments

measure

The argument must be a measure. Null values are omitted from the results. Literal values don't work. The argument must be a field.

group-by level

(Optional) Specifies the level to group the aggregation by. The level added can be any dimension or dimensions independent of the dimensions added to the visual.

The argument must be a dimension field. The group-by level must be enclosed in square brackets []. For more information, see [Level-aware calculation - aggregate \(LAC-A\) functions](#).

Examples

The following example returns the median sales value for each region. It is compared to the total, maximum, and minimum sales.

```
median({Sales})
```

Field wells

Group by
Region

Value
Sales (Sum)
max({Sales}) (Custom)
min({Sales}) (Custom)
median({Sales}) (Custom)

Sheet 1 +

Max({sales}), Sum of Sales, Min({sales}), and Median({sales}) by Region

Region	Sales	max({Sales})	min({Sales})	median({Sales})
AMER	837,849.6	13,999.96	0.99	60.34
APJ	415,464.24	17,499.95	0.44	44.43
EMEA	1,043,887.02	22,638.48	0.85	53.42

You can also specify at what level to group the computation using one or more dimensions in the view or in your dataset. This is called a LAC-A function. For more information about LAC-A functions, see [Level-aware calculation - aggregate \(LAC-A\) functions](#). The following example

calculates the median sales at the Country level, but not across other dimensions (Region) in the visual.

```
median({Sales}, [Country])
```

Median(sales), Sum of Median(sales and [Country]) by Region

Regi...	median(Sales)	median(Sales, [Country])
AMER	60.34	489.27
APJ	44.43	641.82
EMEA	53.42	3,656.93

medianIf

Based on a conditional statement, the medianIf aggregation returns the median value of the specified measure, grouped by the chosen dimension or dimensions. For example, `medianIf(Revenue, SaleDate >= ${BasePeriodStartDate} AND SaleDate <= ${BasePeriodEndDate})` returns the median revenue grouped by the (optional) chosen dimension, if the condition evaluates to true.

Syntax

```
medianIf(measure, condition)
```

Arguments

measure

The argument must be a measure. Null values are omitted from the results. Literal values don't work. The argument must be a field.

condition

One or more conditions in a single statement.

min

The min function returns the minimum value of the specified measure or date, grouped by the chosen dimension or dimensions. For example, `min(return_rate)` returns the minimum rate of returns grouped by the (optional) chosen dimension.

Syntax

```
min(measure, [group-by level])
```

Arguments

measure

The argument must be a measure or a date. Null values are omitted from the results. Literal values don't work. The argument must be a field.

Minimum dates work only in the **Value** field well of tables and pivot tables.

group-by level

(Optional) Specifies the level to group the aggregation by. The level added can be any dimension or dimensions independent of the dimensions added to the visual.

The argument must be a dimension field. The group-by level must be enclosed in square brackets []. For more information, see [Level-aware calculation - aggregate \(LAC-A\) functions](#).

Examples

The following example returns the minimum sales value for each region. It is compared to the total, max, and median sales.

```
min({Sales})
```

Field wells

Group by

Region

Value

Sales (Sum)

max({Sales}) (Custom)

min({Sales}) (Custom)

median({Sales}) (Custom)

Sheet 1

+

Max({sales}), Sum of Sales, Min({sales}), and Median({sales}) by Region

Region	Sales	max({Sales})	min({Sales})	median({Sales})
AMER	837,849.6	13,999.96	0.99	60.34
APJ	415,464.24	17,499.95	0.44	44.43
EMEA	1,043,887.02	22,638.48	0.85	53.42

You can also specify at what level to group the computation using one or more dimensions in the view or in your dataset. This is called a LAC-A function. For more information about LAC-A functions, see [Level-aware calculation - aggregate \(LAC-A\) functions](#). The following example calculates the minimum sales at the Country level, but not across other dimensions (Region) in the visual.

min({Sales}, [Country])

Min(sales), Sum of Min(sales and [Country]) by Region

Regi...	min(Sales)	min(Sales, [Country])
AMER	0.99	22.65
APJ	0.44	28.18
EMEA	0.85	1,716.48

minIf

Based on a conditional statement, the `minIf` function returns the minimum value of the specified measure, grouped by the chosen dimension or dimensions. For example, `minIf(ProdRev, CalendarDay >= ${BasePeriodStartDate} AND CalendarDay <= ${BasePeriodEndDate} AND SourcingType <> 'Indirect')` returns the minimum rate of returns grouped by the (optional) chosen dimension, if the condition evaluates to true.

Syntax

```
minIf(measure, condition)
```

Arguments

measure

The argument must be a measure. Null values are omitted from the results. Literal values don't work. The argument must be a field.

condition

One or more conditions in a single statement.

percentile

The `percentile` function calculates the percentile of the values in measure, grouped by the dimension that's in the field well. There are two varieties of percentile calculation available in Quick:

- [percentileCont](#) uses linear interpolation to determine result.
- [percentileDisc \(percentile\)](#) uses actual values to determine result.

The `percentile` function is an alias of `percentileDisc`.

percentileCont

The `percentileCont` function calculates percentile based on a continuous distribution of the numbers in the measure. It uses the grouping and sorting that are applied in the field wells. It answers questions like: What values are representative of this percentile? To return an exact

percentile value that might not be present in your dataset, use `percentileCont`. To return the nearest percentile value that is present in your dataset, use `percentileDisc` instead.

Syntax

```
percentileCont(expression, percentile, [group-by level])
```

Arguments

measure

Specifies a numeric value to use to compute the percentile. The argument must be a measure or metric. Nulls are ignored in the calculation.

percentile

The percentile value can be any numeric constant 0–100. A percentile value of 50 computes the median value of the measure.

group-by level

(Optional) Specifies the level to group the aggregation by. The level added can be any dimension or dimensions independent of the dimensions added to the visual.

The argument must be a dimension field. The group-by level must be enclosed in square brackets []. For more information, see [Level-aware calculation - aggregate \(LAC-A\) functions](#).

Returns

The result of the function is a number.

Usage notes

The `percentileCont` function calculates a result based on a continuous distribution of the values from a specified measure. The result is computed by linear interpolation between the values after ordering them based on settings in the visual. It's different from `percentileDisc`, which simply returns a value from the set of values that are aggregated over. The result from `percentileCont` might or might not exist in the values from the specified measure.

Examples of percentileCont

The following examples help explain how `percentileCont` works.

Example Comparing median, percentileCont, and percentileDisc

The following example shows the median for a dimension (category) by using the median, percentileCont, and percentileDisc functions. The median value is the same as the percentileCont value. percentileCont interpolates a value, which might or might not be in the data set. However, because percentileDisc always displays a value that exists in the dataset, the two results might not match. The last column in this example shows the difference between the two values. The code for each calculated field is as follows:

- `50%Cont = percentileCont( example , 50 )`
- `median = median( example )`
- `50%Disc = percentileDisc( example , 50 )`
- `Cont-Disc = percentileCont( example , 50 ) - percentileDisc( example , 50 )`
- `example = left( category, 1 )` (To make a simpler example, we used this expression to shorten the names of categories down to their first letter.)

example	median	50%Cont	50%Disc	Cont-Disc
A	22.48	22.48	22.24	0.24
B	20.96	20.96	20.95	0.01
C	24.92	24.92	24.92	0
D	24.935	24.935	24.92	0.015
E	14.48	14.48	13.99	0.49

Example 100th percentile as maximum

The following example shows a variety of percentileCont values for the example field. The calculated fields `n%Cont` are defined as `percentileCont( {example} , n )`. The interpolated values in each column represent the numbers that fall into that percentile bucket. In some cases, the actual data values match the interpolated values. For example, the column `100%Cont` shows the same value for every row because 6783.02 is the highest number.

example	50%Cont	75%Cont	99%Cont	100%Cont
A	20.97	84.307	699.99	6783.02
B	20.99	88.84	880.98	6783.02

C	20.99	90.48	842.925	6783.02
D	21.38	85.99	808.49	6783.02

You can also specify at what level to group the computation using one or more dimensions in the view or in your dataset. This is called a LAC-A function. For more information about LAC-A functions, see [Level-aware calculation - aggregate \(LAC-A\) functions](#). The following example calculates the 30th percentile based on a continuous distribution of the numbers at the Country level, but not across other dimensions (Region) in the visual.

```
percentileCont({Sales}, 30, [Country])
```

Percentilecont(sales, 30), Sum of Percentilecont(sales, 30 and [Country]) by Region

Regi...	percentileCont(Sales, 30)	percentileCont(Sales, 30, [Country])
AMER	23.84	188.74
APJ	16.77	239.59
EMEA	22.32	2,480.17

percentileDisc (percentile)

The percentileDisc function calculates the percentile based on the actual numbers in measure. It uses the grouping and sorting that are applied in the field wells. The percentile function is an alias of percentileDisc.

Use this function to answer the following question: Which actual data points are present in this percentile? To return the nearest percentile value that is present in your dataset, use percentileDisc. To return an exact percentile value that might not be present in your dataset, use percentileCont instead.

Syntax

```
percentileDisc(expression, percentile, [group-by level])
```

Arguments

measure

Specifies a numeric value to use to compute the percentile. The argument must be a measure or metric. Nulls are ignored in the calculation.

percentile

The percentile value can be any numeric constant 0–100. A percentile value of 50 computes the median value of the measure.

group-by level

(Optional) Specifies the level to group the aggregation by. The level added can be any dimension or dimensions independent of the dimensions added to the visual.

The argument must be a dimension field. The group-by level must be enclosed in square brackets []. For more information, see [Level-aware calculation - aggregate \(LAC-A\) functions](#).

Returns

The result of the function is a number.

Usage notes

`percentileDisc` is an inverse distribution function that assumes a discrete distribution model. It takes a percentile value and a sort specification and returns an element from the given set.

For a given percentile value *P*, `percentileDisc` uses the sorted values in the visual and returns the value with the smallest cumulative distribution value that is greater than or equal to *P*.

Examples of `percentileDisc`

The following examples help explain how `percentileDisc` works.

Example Comparing median, `percentileDisc`, and `percentileCont`

The following example shows the median for a dimension (category) by using the `percentileCont`, and `percentileDisc`, and `median` functions. The median value is the same as the `percentileCont` value. `percentileCont` interpolates a value, which might or might not be in the data set. However, because `percentileDisc` always displays the closest value that exists in the dataset, the two results might not match. The last column in this example shows the difference between the two values. The code for each calculated field is as follows:

- `50%Cont = percentileCont( example , 50 )`
- `median = median( example )`

- `50%Disc = percentileDisc( example , 50 )`
- `Cont-Disc = percentileCont( example , 50 ) - percentileDisc( example , 50 )`
- `example = left( category, 1 )` (To make a simpler example, we used this expression to shorten the names of categories down to their first letter.)

example	median	50%Cont	50%Disc	Cont-Disc
A	22.48	22.48	22.24	0.24
B	20.96	20.96	20.95	0.01
C	24.92	24.92	24.92	0
D	24.935	24.935	24.92	0.015
E	14.48	14.48	13.99	0.49

Example 100th percentile as maximum

The following example shows a variety of `percentileDisc` values for the `example` field. The calculated fields `n%Disc` are defined as `percentileDisc( {example} , n)`. The values in each column are actual numbers that come from the dataset.

example	50%Disc	75%Disc	99%Disc	100%Disc
A	20.97	73.98	699.99	6783.02
B	42.19	88.84	820.08	6783.02
C	30.52	90.48	733.44	6783.02
D	41.38	85.99	901.29	6783.0

You can also specify at what level to group the computation using one or more dimensions in the view or in your dataset. This is called a LAC-A function. For more information about LAC-A functions, see [Level-aware calculation - aggregate \(LAC-A\) functions](#). The following example calculates the 30th percentile based on a continuous distribution of the numbers at the Country level, but not across other dimensions (Region) in the visual.

```
percentile({Sales}, 30, [Country])
```

Percentile(sales, 30), Sum of Percentile(sales, 30 and [Country]) by Region

Regi...	percentile(Sales, 30)	percentile(Sales, 30, [Country])
AMER	23.84	187.64
APJ	16.74	237.73
EMEA	22.32	2,472.33

periodToDateAvg

The `periodToDateAvg` function averages the set of numbers in the specified measure for a given time granularity (for instance, a quarter) up to a point in time, relative to that period.

Syntax

```
periodToDateAvg(  
  measure,  
  dateTime,  
  period,  
  endDate (optional))
```

Arguments

measure

The argument must be a field. Null values are omitted from the results. Literal values don't work.

dateTime

The Date dimension over which you're computing PeriodToDate aggregations.

period

The time period across which you're computing the computation. Granularity of YEAR means YearToDate computation, Quarter means QuarterToDate, and so on. Valid granularities include YEAR, QUARTER, MONTH, WEEK, DAY, HOUR, MINUTE, and SECONDS.

endDate

(Optional) The date dimension that you're ending computing periodToDate aggregations. It defaults to `now()` if omitted.

Example

The following example calculates the week-to-date minimum fare amount per payment type, for the week of 06-30-21. For simplicity in the example, we filtered out only a single payment. 06-30-21 is Wednesday. Quick begins the week on Sundays. In our example, that is 06-27-21.

```
periodToDateAvg(fare_amount, pickupDatetime, WEEK, parseDate("06-30-2021", "MM-dd-yyyy"))
```

Field wells

Group by

payment_type

Value

PTDAverage (Custom)

Sheet 1

+

Ptdaverage by Payment_type

payment_ty...	PTDAverage
1	13.74

(A+B+C)/3

Average of Fare_amount by Pickupdatetime and Payment_type

pickupDatetime	payment_type	fare_amount
Jun 30, 2021	1	14.33
Jun 29, 2021	1	C 13.93
Jun 28, 2021	1	B 13.44
Jun 27, 2021	1	A 13.85
Jun 26, 2021	1	14.81

periodToDateCount

The periodToDateCount function calculates the number of values in a dimension or measure, including duplicates, for a given time granularity (for instance, a quarter) up to a point in time, relative to that period.

Syntax

```
periodToDateCount(  
    measure,  
    dateTime,  
    period,  
    endDate (optional))
```

Arguments

measure

The argument must be a field. Null values are omitted from the results. Literal values don't work.

dateTime

The Date dimension over which you're computing PeriodToDate aggregations.

period

The time period across which you're computing the computation. Granularity of YEAR means YearToDate computation, Quarter means QuarterToDate, and so on. Valid granularities include YEAR, QUARTER, MONTH, WEEK, DAY, HOUR, MINUTE, and SECONDS.

endDate

(Optional) The date dimension that you're ending computing periodToDate aggregations. It defaults to now() if omitted.

Example

The following example calculates the week-to-date minimum fare amount per payment type, for the week of 06-30-21. For simplicity in the example, we filtered out only a single payment. 06-30-21 is Wednesday. Quick begins the week on Sundays. In our example, that is 06-27-21.

```
periodToDateCount(fare_amount, pickUpDatetime, WEEK, parseDate("06-30-2021", "MM-dd-yyyy"))
```


Field wells

Group by

payment_type

Value

PTDCount (Custom)

Sheet 1



Ptdcount by Payment_type

payment_ty...	PTDCount
1	701,602

A+B+C

Count of Fare_amount by Pickupdatetime and Payment_type

pickupDatetime	payment_type	fare_amount
Jun 30, 2021	1	240,926
Jun 29, 2021	1	C 240,644
Jun 28, 2021	1	B 230,231
Jun 27, 2021	1	A 230,727
Jun 26, 2021	1	209,895

periodToDateMax

The `periodToDateMax` function returns the maximum value of the specified measure for a given time granularity (for instance, a quarter) up to a point in time, relative to that point.

Syntax

```
periodToDateMax(
  measure,
  dateTime,
  period,
  endDate (optional))
```

Arguments

measure

The argument must be a field. Null values are omitted from the results. Literal values don't work.

dateTime

The Date dimension over which you're computing PeriodToDate aggregations.

period

The time period across which you're computing the computation. Granularity of YEAR means YearToDate computation, Quarter means QuarterToDate, and so on. Valid granularities include YEAR, QUARTER, MONTH, WEEK, DAY, HOUR, MINUTE, and SECONDS.

endDate

(Optional) The date dimension that you're ending computing periodToDate aggregations. It defaults to now() if omitted.

Example

The following example calculates the week-to-date minimum fare amount per payment type, for the week of 06-30-21. For simplicity in the example, we filtered out only a single payment. 06-30-21 is Wednesday. Quick begins the week on Sundays. In our example, that is 06-27-21.

```
periodToDateMax(fare_amount, pickUpDatetime, WEEK, parseDate("06-30-2021", "MM-dd-yyyy"))
```

Field wells

Group by
payment_type

Value
PTDMax (Custom)

Sheet 1 +

Ptdmax by Payment_type

payment_ty...	PTDMax
1	490

Max of Fare_amount by Pickupdatetime and Payment_type

pickupDatetime	payment_type	fare_amount
Jun 30, 2021	1	500
Jun 29, 2021	1	400
Jun 28, 2021	1	320
Jun 27, 2021	1	490
Jun 26, 2021	1	500

periodToDateMedian

The `periodToDateMedian` function returns the median value of the specified measure for a given time granularity (for instance, a quarter) up to a point in time, relative to that period.

Syntax

```
periodToDateMedian(
  measure,
  dateTime,
  period,
  endDate (optional))
```

Arguments

measure

The argument must be a field. Null values are omitted from the results. Literal values don't work.

dateTime

The Date dimension over which you're computing PeriodToDate aggregations.

period

The time period across which you're computing the computation. Granularity of YEAR means YearToDate computation, Quarter means QuarterToDate, and so on. Valid granularities include YEAR, QUARTER, MONTH, WEEK, DAY, HOUR, MINUTE, and SECONDS.

endDate

(Optional) The date dimension that you're ending computing periodToDate aggregations. It defaults to now() if omitted.

Example

The following example calculates the week-to-date minimum fare amount per payment type, for the week of 06-30-21. For simplicity in the example, we filtered out only a single payment. 06-30-21 is Wednesday. Quick begins the week on Sundays. In our example, that is 06-27-21.

```
periodToDateMedian(fare_amount, pickUpDatetime, WEEK, parseDate("06-30-2021", "MM-dd-yyyy"))
```

Field wells

Group by

payment_type

Value

PTDMedian (Custom)

Sheet 1 ▾ +

Ptdmedian by Payment_type

payment_ty...	PTDMedian
1	10

Median of Fare_amount by Pickupdatetime and Payment_type

pickupDatetime	payment_type	fare_amount
Jun 30, 2021	1	10.50
Jun 29, 2021	1	10.00
Jun 28, 2021	1	10.00
Jun 27, 2021	1	10.00
Jun 26, 2021	1	10.50

periodToDateMin

The `periodToDateMin` function returns the minimum value of the specified measure or date, or a given time granularity (for instance, a quarter) up to a point in time, relative to that period.

Syntax

```
periodToDateMin(
  measure,
  dateTime,
  period,
  endDate (optional))
```

Arguments

measure

The argument must be a field. Null values are omitted from the results. Literal values don't work.

dateTime

The Date dimension over which you're computing PeriodToDate aggregations.

period

The time period across which you're computing the computation. Granularity of YEAR means YearToDate computation, Quarter means QuarterToDate, and so on. Valid granularities include YEAR, QUARTER, MONTH, WEEK, DAY, HOUR, MINUTE, and SECONDS.

endDate

(Optional) The date dimension that you're ending computing periodToDate aggregations. It defaults to now() if omitted.

Example

The following example calculates the week-to-date minimum fare amount per payment type, for the week of 06-30-21. For simplicity in the example, we filtered out only a single payment. 06-30-21 is Wednesday. Quick begins the week on Sundays. In our example, that is 06-27-21.

```
periodToDateMin(fare_amount, pickUpDatetime, WEEK, parseDate("06-30-2021", "MM-dd-yyyy"))
```

Field wells

Group by
payment_type

Value
PTDMin (Custom)

Sheet 1 +

Ptdmin by Payment_type

payment_ty...	PTDMin
1	-7

Min of Fare_amount by Pickupdatetime and Payment_type

pickupDatetime	payment_type	fare_amount
Jun 30, 2021	1	-52
Jun 29, 2021	1	0
Jun 28, 2021	1	0
Jun 27, 2021	1	-7
Jun 26, 2021	1	0

periodToDatePercentile

The `periodToDatePercentile` function calculates the percentile based on the actual numbers in measure for a given time granularity (for instance, a quarter) up to a point in time, relative to that period. It uses the grouping and sorting that are applied in the field wells.

To return the nearest percentile value that is present in your dataset, use `periodToDatePercentile`. To return an exact percentile value that might not be present in your dataset, use `periodToDatePercentileCont` instead.

Syntax

```
periodToDatePercentile(
  measure,
```

```
percentile,  
dateTime,  
period,  
endDate (optional))
```

Arguments

measure

The argument must be a field. Null values are omitted from the results. Literal values don't work.

percentile

The percentile value can be any numeric constant 0-100. A percentile of 50 computes the median value of the measure.

dateTime

The Date dimension over which you're computing PeriodToDate aggregations.

period

The time period across which you're computing the computation. Granularity of YEAR means YearToDate computation, Quarter means QuarterToDate, and so on. Valid granularities include YEAR, QUARTER, MONTH, WEEK, DAY, HOUR, MINUTE, and SECONDS.

endDate

(Optional) The date dimension that you're ending computing periodToDate aggregations. It defaults to now() if omitted.

Example

The following example calculates the week-to-date, 90th percentile of fare amount per payment type for the week of 06-30-21. For simplicity in the example, we filtered out only a single payment. 06-30-21 is Wednesday. Quick begins the week on Sundays. In our example. that is 06-27-21.

```
periodToDatePercentile(fare_amount, 90, pickupDatetime, WEEK, parseDate("06-30-2021",  
"MM-dd-yyyy"))
```


Field wells

Group by
payment_type

Value
PTDPercentile (Custom)

Sheet 1 +

Ptdpercentile by Payment_type

payment_ty...	PTDPercentile
1	27

P90 of Fare_amount by Pickupdatetime and Payment_type

pickupDatetime	payment_type	fare_amount
Jun 30, 2021	1	29.00
Jun 29, 2021	1	27.50
Jun 28, 2021	1	26.00
Jun 27, 2021	1	28.50
Jun 26, 2021	1	30.00

periodToDatePercentileCont

The `periodToDatePercentileCont` function calculates percentile based on a continuous distribution of the numbers in the measure for a given time granularity (for instance, a quarter) up to a point in time in that period. It uses the grouping and sorting that are applied in the field wells.

To return an exact percentile value that might not be present in your dataset, use `periodToDatePercentileCont`. To return the nearest percentile value that is present in your dataset, use `periodToDatePercentile` instead.

Syntax

```
periodToDatePercentileCont(
  measure,
```

```
percentile,  
dateTime,  
period,  
endDate (optional))
```

Arguments

measure

The argument must be a field. Null values are omitted from the results. Literal values don't work.

percentile

The percentile value can be any numeric constant 0-100. A percentile of 50 computes the median value of the measure.

dateTime

The Date dimension over which you're computing PeriodToDate aggregations.

period

The time period across which you're computing the computation. Granularity of YEAR means YearToDate computation, Quarter means QuarterToDate, and so on. Valid granularities include YEAR, QUARTER, MONTH, WEEK, DAY, HOUR, MINUTE, and SECONDS.

endDate

(Optional) The date dimension that you're ending computing periodToDate aggregations. It defaults to now() if omitted.

Example

The following example calculates the week-to-date, 90th percentile of fare amount per payment type for the week of 06-30-21. For simplicity in the example, we filtered out only a single payment. 06-30-21 is Wednesday. Quick begins the week on Sundays. In our example, that is 06-27-21.

```
periodToDatePercentileCont(fare_amount, 90, pickupDatetime, WEEK,  
parseDate("06-30-2021", "MM-dd-yyyy"))
```

Field wells

Group by
payment_type

Value
PTDPercentileCont (Custom)

Sheet 1 +

Ptdpercentilecont by Payment_type

payment_ty...	PTDPercentileCont
1	27

P90 of Fare_amount by Pickupdatetime and Payment_type

pickupDatetime	payment_type	fare_amount
Jun 30, 2021	1	29.00
Jun 29, 2021	1	27.50
Jun 28, 2021	1	26.00
Jun 27, 2021	1	28.50
Jun 26, 2021	1	30.00

periodToDateStDev

The `periodToDateStDev` function calculates the standard deviation of the set of numbers in the specified measure for a given time granularity (for instance, a quarter) up to a point in time, based on a sample and relative to that period.

Syntax

```
periodToDateStDev(
  measure,
  dateTime,
  period,
  endDate (optional))
```

Arguments

measure

The argument must be a field. Null values are omitted from the results. Literal values don't work.

dateTime

The Date dimension over which you're computing PeriodToDate aggregations.

period

(Optional) The time period across which you're computing the computation. Granularity of YEAR means YearToDate computation, Quarter means QuarterToDate, and so on. Valid granularities include YEAR, QUARTER, MONTH, WEEK, DAY, HOUR, MINUTE, and SECONDS.

endDate

(Optional) The date dimension that you're ending computing periodToDate aggregations. It defaults to now() if omitted.

Example

The following example calculates the week-to-date minimum fare amount per payment type, for the week of 06-30-21. For simplicity in the example, we filtered out only a single payment. 06-30-21 is Wednesday. Quick begins the week on Sundays. In our example, that is 06-27-21.

```
periodToDateStDev(fare_amount, pickUpDatetime, WEEK, parseDate("06-30-2021", "MM-dd-yyyy"))
```

Field wells

Group by
payment_type

Value
PTDStDev (Custom)

Sheet 1 +

Ptdstdev by Payment_type

payment_ty...	PTDStDev
1	11.78

Standard deviation of Fare_amount by Pickupdatetime and Payment_type

pickupDatetime	payment_type	fare_amount
Jun 30, 2021	1	12.26
Jun 29, 2021	1	11.65
Jun 28, 2021	1	11.45
Jun 27, 2021	1	12.21
Jun 26, 2021	1	12.87

periodToDateStDevP

The `periodToDateStDevP` function calculates the population standard deviation of the set of numbers in the specified measure, for a given time granularity (for instance, a quarter) up to a point in time, based on a sample in that period.

Syntax

```
periodToDateStDevP(
  measure,
  dateTime,
  period,
  endDate (optional))
```

Arguments

measure

The argument must be a field. Null values are omitted from the results. Literal values don't work.

dateTime

The Date dimension over which you're computing PeriodToDate aggregations.

period

The time period across which you're computing the computation. Granularity of YEAR means YearToDate computation, Quarter means QuarterToDate, and so on. Valid granularities include YEAR, QUARTER, MONTH, WEEK, DAY, HOUR, MINUTE, and SECONDS.

endDate

(Optional) The date dimension that you're ending computing periodToDate aggregations. It defaults to now() if omitted.

Example

The following example calculates the week-to-date minimum fare amount per payment type, for the week of 06-30-21. For simplicity in the example, we filtered out only a single payment. 06-30-21 is Wednesday. Quick begins the week on Sundays. In our example, that is 06-27-21.

```
periodToDateStDevP(fare_amount, pickUpDatetime, WEEK, parseDate("06-30-2021", "MM-dd-yyyy"))
```

Field wells

Group by

payment_type

Value

PTDStDevP (Custom)

Sheet 1



Ptdstdevp by Payment_type

payment_ty...	PTDStDevP
1	11.78

Standard deviation - population of Fare_amount by Pickupdatetime and Payment_type

pickupDatetime	payment_type	fare_amount
Jun 30, 2021	1	12.26
Jun 29, 2021	1	11.65
Jun 28, 2021	1	11.45
Jun 27, 2021	1	12.21
Jun 26, 2021	1	12.87

periodToDateSum

The `periodToDateSum` function adds the specified measure for a given time granularity (for instance, a quarter) up to a point in time, relative to that period.

Syntax

```
periodToDateSum(
  measure,
  dateTime,
  period,
  endDate)
```

Arguments

measure

The argument must be a field. Null values are omitted from the results. Literal values don't work.

dateTime

The Date dimension over which you're computing PeriodToDate aggregations.

period

The time period across which you're computing the computation. Granularity of YEAR means YearToDate computation, Quarter means QuarterToDate, and so on. Valid granularities include YEAR, QUARTER, MONTH, WEEK, DAY, HOUR, MINUTE, and SECONDS.

endDate

(Optional) The date dimension that you're ending computing periodToDate aggregations. It defaults to now() if omitted.

Example

The following function calculates the week to date sum of fare amount per payment, for the week of 06-30-21. For simplicity in the example, we filtered out only a single payment. 06-30-21 is Wednesday. Quick begins the week on Sundays. In our example, that is 06-27-21.

```
periodToDateSum(fare_amount, pickUpDateTime, WEEK, parseDate("06-30-2021", "MM-dd-yyyy"))
```


Field wells

Group by
payment_type

Value
PTDSum (Custom)

Sheet 1 +

Ptdsum by Payment_type

payment_ty...	PTDSum
1	9,642,352.38

A+B+C

Sum of Fare_amount by Pickupdatetime and Payment_type

pickupDatetime	payment_type	fare_amount
Jun 30, 2021	1	3,452,526.24
Jun 29, 2021	1	C 3,352,298.42
Jun 28, 2021	1	B 3,095,150.7
Jun 27, 2021	1	A 3,194,903.26
Jun 26, 2021	1	3,108,392.02

periodToDateVar

The `periodToDateVar` function calculates the sample variance of the set of numbers in the specified measure for a given time granularity (for instance, a quarter) up to a point in time in that period.

Syntax

```
periodToDateVar(
  measure,
  dateTime,
  period,
  endDate (optional))
```

Arguments

measure

The argument must be a field. Null values are omitted from the results. Literal values don't work.

dateTime

The Date dimension over which you're computing PeriodToDate aggregations.

period

The time period across which you're computing the computation. Granularity of YEAR means YearToDate computation, Quarter means QuarterToDate, and so on. Valid granularities include YEAR, QUARTER, MONTH, WEEK, DAY, HOUR, MINUTE, and SECONDS.

endDate

(Optional) The date dimension that you're ending computing periodToDate aggregations. It defaults to now() if omitted.

Example

The following example calculates the week-to-date minimum fare amount per payment type, for the week of 06-30-21. For simplicity in the example, we filtered out only a single payment. 06-30-21 is Wednesday. Quick begins the week on Sundays. In our example, that is 06-27-21.

```
periodToDateVar(fare_amount, pickUpDatetime, WEEK, parseDate("06-30-2021", "MM-dd-yyyy"))
```

Field wells

Group by
payment_type

Value
PTDVar (Custom)

Sheet 1 +

Ptdvar by Payment_type

payment_ty...	PTDVar
1	138.65

Variance of Fare_amount by Pickupdatetime and Payment_type

pickupDatetime	payment_type	fare_amount
Jun 30, 2021	1	150.22
Jun 29, 2021	1	135.78
Jun 28, 2021	1	131.13
Jun 27, 2021	1	149.02
Jun 26, 2021	1	165.76

periodToDateVarP

The `periodToDateVarP` function calculates the population variance of the set of numbers in the specified measure for a given time granularity (for instance, a quarter) up to a point in time, relevant to that period.

Syntax

```
periodToDateVarP(
  measure,
  dateTime,
  period,
  endDate (optional))
```

Arguments

measure

The argument must be a field. Null values are omitted from the results. Literal values don't work.

dateTime

The Date dimension over which you're computing PeriodToDate aggregations.

period

The time period across which you're computing the computation. Granularity of YEAR means YearToDate computation, Quarter means QuarterToDate, and so on. Valid granularities include YEAR, QUARTER, MONTH, WEEK, DAY, HOUR, MINUTE, and SECONDS.

endDate

(Optional) The date dimension that you're ending computing periodToDate aggregations. It defaults to now() if omitted.

Example

The following example calculates the week-to-date minimum fare amount per payment type, for the week of 06-30-21. For simplicity in the example, we filtered out only a single payment. 06-30-21 is Wednesday. Quick begins the week on Sundays. In our example, that is 06-27-21.

```
periodToDateVarP(fare_amount, pickUpDatetime, WEEK, parseDate("06-30-2021", "MM-dd-yyyy"))
```

Field wells

Group by
payment_type

Value
PTDVarP (Custom)

Sheet 1 +

Ptdvarp by Payment_type

payment_ty...	PTDVarP
1	138.65

Variance - population of Fare_amount by Pickupdatetime and Payment_type

pickupDatetime	payment_type	fare_amount
Jun 30, 2021	1	150.22
Jun 29, 2021	1	135.78
Jun 28, 2021	1	131.13
Jun 27, 2021	1	149.02
Jun 26, 2021	1	165.76

stdev

The `stdev` function calculates the standard deviation of the set of numbers in the specified measure, grouped by the chosen dimension or dimensions, based on a sample.

Syntax

```
stdev(measure, [group-by level])
```

Arguments

measure

The argument must be a measure. Null values are omitted from the results. Literal values don't work. The argument must be a field.

group-by level

(Optional) Specifies the level to group the aggregation by. The level added can be any dimension or dimensions independent of the dimensions added to the visual.

The argument must be a dimension field. The group-by level must be enclosed in square brackets []. For more information, see [Level-aware calculation - aggregate \(LAC-A\) functions](#).

Examples

The following example returns the standard deviation of test scores for a class, using a sample of the test scores recorded.

```
stdev({Score})
```

You can also specify at what level to group the computation using one or more dimensions in the view or in your dataset. This is called a LAC-A function. For more information about LAC-A functions, see [Level-aware calculation - aggregate \(LAC-A\) functions](#). The following example calculates the standard deviation of test scores at the subject level, but not across other dimensions (Class) in the visual.

```
stdev({Score}, [Subject])
```

stdevp

The `stdevp` function calculates the population standard deviation of the set of numbers in the specified measure, grouped by the chosen dimension or dimensions.

Syntax

```
stdevp(measure, [group-by level])
```

Arguments

measure

The argument must be a measure. Null values are omitted from the results. Literal values don't work. The argument must be a field.

group-by level

(Optional) Specifies the level to group the aggregation by. The level added can be any dimension or dimensions independent of the dimensions added to the visual.

The argument must be a dimension field. The group-by level must be enclosed in square brackets []. For more information, see [Level-aware calculation - aggregate \(LAC-A\) functions](#).

Examples

The following example returns the standard deviation of test scores for a class using all the scores recorded.

```
stdevp({Score})
```

You can also specify at what level to group the computation using one or more dimensions in the view or in your dataset. This is called a LAC-A function. For more information about LAC-A functions, see [Level-aware calculation - aggregate \(LAC-A\) functions](#). The following example calculates the standard deviation of test scores at the subject level, but not across other dimensions (Class) in the visual using all the scores recorded.

```
stdevp({Score}, [Subject])
```

stdevIf

Based on a conditional statement, the `stdevIf` function calculates the standard deviation of the set of numbers in the specified measure, grouped by the chosen dimension or dimensions, based on a sample.

Syntax

```
stdevIf(measure, conditions)
```

Arguments

measure

The argument must be a measure. Null values are omitted from the results. Literal values don't work. The argument must be a field.

condition

One or more conditions in a single statement.

stdevpIf

Based on a conditional statement, the `stdevpIf` function calculates the standard deviation of the set of numbers in the specified measure, grouped by the chosen dimension or dimensions, based on a biased population.

Syntax

```
stdevpIf(measure, conditions)
```

Arguments

measure

The argument must be a measure. Null values are omitted from the results. Literal values don't work. The argument must be a field.

condition

One or more conditions in a single statement.

sum

The `sum` function adds the set of numbers in the specified measure, grouped by the chosen dimension or dimensions. For example, `sum(profit amount)` returns the total profit amount grouped by the (optional) chosen dimension.

Syntax

```
sum(measure, [group-by level])
```

Arguments

measure

The argument must be a measure. Null values are omitted from the results. Literal values don't work. The argument must be a field.

group-by level

(Optional) Specifies the level to group the aggregation by. The level added can be any dimension or dimensions independent of the dimensions added to the visual.

The argument must be a dimension field. The group-by level must be enclosed in square brackets []. For more information, see [Level-aware calculation - aggregate \(LAC-A\) functions](#).

Examples

The following example returns the sum of sales.

```
sum({Sales})
```

You can also specify at what level to group the computation using one or more dimensions in the view or in your dataset. This is called a LAC-A function. For more information about LAC-A functions, see [Level-aware calculation - aggregate \(LAC-A\) functions](#). The following example calculates the sum of sales at the Country level, but not across other dimensions (Region and Product) in the visual.

```
sum(Sales, [Country])
```

Sum(sales), Sum of Sum(sales and [Country]) by Region, Product, and Country

Regi...	Country	Product	sum(Sales)	sum(Sales, [Country])
AMER	Argentina	Big Ol Database	9,899.85	35,764.31
AMER	Argentina	ChatBot Plugin	742.8	35,764.31
AMER	Argentina	ContactMatcher	3,947.81	35,764.31
AMER	Argentina	Data Smasher	1,023.56	35,764.31
AMER	Argentina	FinanceHub	2,728.24	35,764.31
AMER	Argentina	Marketing Suite	2,275.88	35,764.31
AMER	Argentina	Marketing Suite - Gold	4,669.08	35,764.31
AMER	Argentina	OneView	4,204.36	35,764.31
AMER	Argentina	SaaS Connector Pack	950.97	35,764.31
AMER	Argentina	SaaS Connector Pack - Gold	153.7	35,764.31
AMER	Argentina	Site Analytics	3,577.75	35,764.31
AMER	Argentina	Storage	54.12	35,764.31
AMER	Argentina	Support	1,536.19	35,764.31
APJ	Australia	Alchemy	5,919.89	80,166.1
APJ	Australia	Big Ol Database	3,756.31	80,166.1
APJ	Australia	ChatBot Plugin	930.03	80,166.1

sumIf

Based on a conditional statement, the sumIf function adds the set of numbers in the specified measure, grouped by the chosen dimension or dimensions. For example, `sumIf(ProdRev, CalendarDay >= ${BasePeriodStartDate} AND CalendarDay <= ${BasePeriodEndDate} AND SourcingType <> 'Indirect')` returns the total profit amount grouped by the (optional) chosen dimension, if the condition evaluates to true.

Syntax

```
sumIf(measure, conditions)
```

Arguments*measure*

The argument must be a measure. Null values are omitted from the results. Literal values don't work. The argument must be a field.

condition

One or more conditions in a single statement.

Examples

The following example uses a calculated field with `sumIf` to display the sales amount if Segment is equal to SMB.

```
sumIf(Sales, Segment='SMB')
```

Field wells

Group by

Segment

Value

Sales (Sum)

sumIf (Custom)

Lag SumIf +

Sum of Sales and Sumif by Segment

Segment	Sales	sumIf
Enterprise	429,653	
SMB	1,161,401	1,161,401.35
Strategic	706,146	

The following example uses a calculated field with `sumIf` to display the sales amount if Segment is equal to SMB and Order Date greater than year 2022.

```
sumIf(Sales, Segment='SMB' AND {Order Date} >='2022-01-01')
```

var

The `var` function calculates the sample variance of the set of numbers in the specified measure, grouped by the chosen dimension or dimensions.

Syntax

```
var(measure, [group-by level])
```

Arguments

measure

The argument must be a measure. Null values are omitted from the results. Literal values don't work. The argument must be a field.

group-by level

(Optional) Specifies the level to group the aggregation by. The level added can be any dimension or dimensions independent of the dimensions added to the visual.

The argument must be a dimension field. The group-by level must be enclosed in square brackets []. For more information, see [Level-aware calculation - aggregate \(LAC-A\) functions](#).

Examples

The following example returns the variance of a sample of test scores.

```
var({Scores})
```

You can also specify at what level to group the computation using one or more dimensions in the view or in your dataset. This is called a LAC-A function. For more information about LAC-A functions, see [Level-aware calculation - aggregate \(LAC-A\) functions](#). The following example returns the variance of a sample of test scores at the subject level, but not across other dimensions (Class) in the visual.

```
var({Scores}, [Subject])
```

varIf

Based on a conditional statement, the `varIf` function calculates the variance of the set of numbers in the specified measure, grouped by the chosen dimension or dimensions, based on a sample.

Syntax

```
varIf(measure, conditions)
```

Arguments

measure

The argument must be a measure. Null values are omitted from the results. Literal values don't work. The argument must be a field.

condition

One or more conditions in a single statement.

varp

The `varp` function calculates the population variance of the set of numbers in the specified measure, grouped by the chosen dimension or dimensions.

Syntax

```
varp(measure, [group-by level])
```

Arguments

measure

The argument must be a measure. Null values are omitted from the results. Literal values don't work. The argument must be a field.

group-by level

(Optional) Specifies the level to group the aggregation by. The level added can be any dimension or dimensions independent of the dimensions added to the visual.

The argument must be a dimension field. The group-by level must be enclosed in square brackets []. For more information, see [Level-aware calculation - aggregate \(LAC-A\) functions](#).

Examples

The following example returns the variance of a population of test scores.

```
varp({Scores})
```

You can also specify at what level to group the computation using one or more dimensions in the view or in your dataset. This is called a LAC-A function. For more information about LAC-A functions, see [Level-aware calculation - aggregate \(LAC-A\) functions](#). The following example returns the variance of a population test scores at the subject level, but not across other dimensions (Class) in the visual.

```
varp({Scores}, [Subject])
```

varpIf

Based on a conditional statement, the `varpIf` function calculates the variance of the set of numbers in the specified measure, grouped by the chosen dimension or dimensions, based on a biased population.

Syntax

```
varpIf(measure, conditions)
```

Arguments

measure

The argument must be a measure. Null values are omitted from the results. Literal values don't work. The argument must be a field.

condition

One or more conditions in a single statement.

Table calculation functions

When you are analyzing data in a specific visual, you can apply table calculations to the current set of data to discover how dimensions influence measures or each other. *Visualized data* is your result set based on your current dataset, with all the filters, field selections, and customizations applied. To see exactly what this result set is, you can export your visual to a file. A *table calculation function* performs operations on the data to reveal relationships between fields.

In this section, you can find a list of the functions available in table calculations that you can perform on visualized data in Amazon Quick.

To view a list of functions sorted by category, with brief definitions, see [Functions by category](#).

Topics

- [difference](#)
- [distinctCountOver](#)
- [lag](#)
- [lead](#)
- [percentDifference](#)
- [avgOver](#)
- [countOver](#)
- [maxOver](#)
- [minOver](#)
- [percentileOver](#)
- [percentileContOver](#)
- [percentileDiscOver](#)
- [percentOfTotal](#)
- [periodOverPeriodDifference](#)
- [periodOverPeriodLastValue](#)
- [periodOverPeriodPercentDifference](#)
- [periodToDateAvgOverTime](#)
- [periodToDateCountOverTime](#)
- [periodToDateMaxOverTime](#)
- [periodToDateMinOverTime](#)
- [periodToDateSumOverTime](#)
- [stdevOver](#)
- [stdevpOver](#)
- [varOver](#)
- [varpOver](#)
- [sumOver](#)
- [denseRank](#)

- [rank](#)
- [percentileRank](#)
- [runningAvg](#)
- [runningCount](#)
- [runningMax](#)
- [runningMin](#)
- [runningSum](#)
- [firstValue](#)
- [lastValue](#)
- [windowAvg](#)
- [windowCount](#)
- [windowMax](#)
- [windowMin](#)
- [windowSum](#)

difference

The difference function calculates the difference between a measure based on one set of partitions and sorts, and a measure based on another.

The difference function is supported for use with analyses based on SPICE and direct query data sets.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
difference
(
    measure
    , [ sortorder_field ASC_or_DESC, ... ]
    , lookup_index,
    , [ partition field, ... ]
)
```


Arguments

measure

An aggregated measure that you want to see the difference for.

sort order field

One or more measures and dimensions that you want to sort the data by, separated by commas. You can specify either ascending (**ASC**) or descending (**DESC**) sort order.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

lookup index

The lookup index can be positive or negative, indicating a following row in the sort (positive) or a previous row in the sort (negative). The lookup index can be 1–2,147,483,647. For the engines MySQL, MariaDB and Aurora MySQL-Compatible Edition, the lookup index is limited to just 1.

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

Example

The following example calculates the difference between of `sum({Billed Amount})`, sorted by `Customer Region` ascending, compared to the next row, and partitioned by `Service Line`.

```
difference(  
    sum( {Billed Amount} ),  
    [{Customer Region} ASC],  
    1,  
    [{Service Line}]  
)
```

The following example calculates the difference between `Billed Amount` compared to the next line, partitioned by (`[{Customer Region}]`). The fields in the table calculation are in the field wells of the visual.

```
difference(
  sum( {Billed Amount} ),
  [{Customer Region} ASC],
  1
)
```

The red highlights show how each amount is added ($a + b = c$) to show the difference between amounts a and c.

Field wells

Group by

Customer Region

Value

difference (Custom)

Billed Amount (Sum)

Sum of Billed Amount and Difference by Customer Region

Customer Region	difference	Billed Amount
APAC		a 8,390,654
EMEA	+b 2,647,510	=c 11,038,164
US	12,509,001	23,547,165

distinctCountOver

The `distinctCountOver` function calculates the distinct count of the operand partitioned by the specified attributes at a specified level. Supported levels are `PRE_FILTER` and `PRE_AGG`. The operand must be unaggregated.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
distinctCountOver
(
  measure or dimension field
```

```
,[ partition_field, ... ]  
,calculation_level  
)
```

Arguments

measure or dimension field

The measure or dimension that you want to do the calculation for, for example {Sales Amt}. Valid values are PRE_FILTER and PRE_AGG.

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

calculation level

(Optional) Specifies the calculation level to use:

- **PRE_FILTER** – Prefilter calculations are computed before the dataset filters.
- **PRE_AGG** – Preaggregate calculations are computed before applying aggregations and top and bottom *N* filters to the visuals.

This value defaults to POST_AGG_FILTER when blank. POST_AGG_FILTER is not a valid level for this operation and will result in an error message. For more information, see [Using level-aware calculations in Amazon Quick](#).

Example

The following example gets the distinct count of Sales partitioned over City and State at the PRE_AGG level.

```
distinctCountOver  
(  
  Sales,  
  [City, State], PRE_AGG  
)
```

lag

The `lag` function calculates the lag (previous) value for a measure based on specified partitions and sorts.

`lag` is supported for use with analyses based on SPICE and direct query data sets.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
lag
(  
  lag  
  (  
    measure  
    , [ sortorder_field ASC_or_DESC, ... ]  
    , lookup_index  
    , [ partition_field, ... ]  
  )  
)
```

Arguments

measure

The measure that you want to get the lag for. This can include an aggregate, for example `sum({Sales Amt})`.

sort order field

One or more measures and dimensions that you want to sort the data by, separated by commas. You can specify either ascending (**ASC**) or descending (**DESC**) sort order.

Each field in the list is enclosed in `{ }` (curly braces), if it is more than one word. The entire list is enclosed in `[ ]` (square brackets).

lookup index

The lookup index can be positive or negative, indicating a following row in the sort (positive) or a previous row in the sort (negative). The lookup index can be 1–2,147,483,647. For the engines MySQL, MariaDB, and Amazon Aurora MySQL-Compatible Edition, the lookup index is limited to just 1.

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

Example

The following example calculates the previous `sum(sales)`, partitioned by the state of origin, in the ascending sort order on `cancellation_code`.

```
lag
(
    sum(Sales),
    [cancellation_code ASC],
    1,
    [origin_state_nm]
)
```

The following example uses a calculated field with `lag` to display sales amount for the previous row next to the amount for the current row, sorted by `Order Date`. The fields in the table calculation are in the field wells of the visual.

```
lag(
    sum({Sales}),
    [{Order Date} ASC],
    1
)
```

The following screenshot shows the results of the example.

Field wells

Group by
Order Date

Value
Sales (Sum)
lag1 (Custom)

Sheet 1 +

Sum of Sales and Lag1 by Order Date

Order Date	Sales	lag1
Jan 4, 2020	16	
Jan 5, 2020	288	16
Jan 6, 2020	20	288
Jan 7, 2020	4,407	20
Jan 8, 2020	87	4,407
Jan 10, 2020	41	87
Jan 11, 2020	55	41
Jan 12, 2020	10	55

The following example uses a calculated field with `lag` to display the sales amount for the previous row next to the amount for the current row, sorted by `Order Date` partitioned by `Segment`.

```
lag
(
  sum(Sales),
  [Order Date ASC],
  1, [Segment]
)
```

The following screenshot shows the results of the example.

Field wells

Group by

Order Date

Segment

Value

Sales (Sum)

lag2 (Custom)

Sheet 1 Sheet 2 +

Sum of Sales and Lag2 by Order Date and Segment

Order Date	Segment	Sales	lag2
Jan 4, 2020	SMB	16	
Jan 5, 2020	Enterprise	288	
Jan 6, 2020	SMB	20	16
Jan 7, 2020	Enterprise	4,375	288
Jan 7, 2020	SMB	19	20
Jan 7, 2020	Strategic	13	

lead

The `lead` function calculates the lead (following) value for a measure based on specified partitions and sorts.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
lead
(
    measure
    , [ sortorder_field ASC_or_DESC, ... ]
    , lookup_index
    , [ partition_field, ... ]
)
```

Arguments

measure

The measure that you want to get the lead for. This can include an aggregate, for example `sum({Sales Amt})`.

sort order field

One or more measures and dimensions that you want to sort the data by, separated by commas. You can specify either ascending (**ASC**) or descending (**DESC**) sort order.

Each field in the list is enclosed in `{ }` (curly braces), if it is more than one word. The entire list is enclosed in `[ ]` (square brackets).

lookup index

The lookup index can be positive or negative, indicating a following row in the sort (positive) or a previous row in the sort (negative). The lookup index can be 1–2,147,483,647. For the engines MySQL, MariaDB, and Amazon Aurora MySQL-Compatible Edition, the lookup index is limited to just 1.

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in `{ }` (curly braces), if it is more than one word. The entire list is enclosed in `[ ]` (square brackets).

Example

The following example calculates the next `sum(sales)`, partitioned by the state of origin, in the ascending sort order on `cancellation_code`.


```
lead
(
    sum(sales),
    [cancellation_code ASC],
    1,
    [origin_state_nm]
)
```


The following example uses a calculated field with lead to display the amount for the next row beside the amount for the current row, sorted by Customer Segment. The fields in the table calculation are in the field wells of the visual.

```
lead(
  sum({Billed Amount}),
  [{Customer Segment} ASC],
  1
)
```

The following screenshot shows the results of the example.

Field wells



percentDifference

The percentDifference function calculates the percentage difference between the current value and a comparison value, based on partitions, sorts, and lookup index.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
percentDifference
(
  measure
```

```
,[ sortorder_field ASC_or_DESC, ... ]  
,lookup index  
,[ partition_field, ... ]  
)
```

Arguments

measure

An aggregated measure that you want to see the percent difference for.

sort order field

One or more measures and dimensions that you want to sort the data by, separated by commas. You can specify either ascending (**ASC**) or descending (**DESC**) sort order.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

lookup index

The lookup index can be positive or negative, indicating a following row in the sort (positive) or a previous row in the sort (negative). The lookup index can be 1–2,147,483,647. For the engines MySQL, MariaDB and Aurora MySQL-Compatible Edition, the lookup index is limited to just 1.

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

Example

The following example calculates the percentage of difference between the sum(Sales) for the current and the previous State, sorted by Sales.

```
percentDifference  
(  
  sum(amount),  
  [sum(amount) ASC],  
  -1,  
  [State]
```

)

The following example calculates the percent that a specific Billed Amount is of another Billed Amount, sorted by ([{Customer Region} ASC]). The fields in the table calculation are in the field wells of the visual.

```
percentDifference
(
  sum( {Billed Amount} ),
  [{Customer Region} ASC],
  1
)
```

The following screenshot shows the results of the example. The red letters show that the total Billed Amount for the Customer Region **APAC** is 24 percent less than the amount for the **EMEA** region.

Field wells

Group by

Customer Region

Value

Billed Amount (Sum)

percentDifference (Custom)

Percentdifference and Sum of Billed Amount by Customer Region		
Customer Region	Billed Amount	percentDifference
APAC	a 8,390,654	c -24.0%
EMEA	b 11,038,164	-53.1%
US	23,547,165	(b-a) / b = c

avgOver

The avgOver function calculates the average of a measure partitioned by a list of dimensions.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
avgOver
(
    measure
    ,[ partition_field, ... ]
    ,calculation level
)
```

The following example shows the average Billed Amount over Customer Region. The fields in the table calculation are in the field wells of the visual.

```
avgOver
(
    sum({Billed Amount}),
    [{Customer Region}]
)
```

The following screenshot shows the results of the example. With the addition of Service Line, the total amount billed for each is displayed, and the average of these three values displays in the calculated field.

Field wells

Group by

Customer Region

Service Line

Value

Billed Amount (Sum)

avgOver (Custom)

Sum of Billed Amount and Avgover by Customer Region and Service Line

Customer Region	Service Line	Billed Amount	avgOver
APAC	Billing	3,569,780	2,796,885
APAC	HR	3,441,106	2,796,885
APAC	Marketing	1,379,768	2,796,885
EMEA	Billing	4,446,586	3,679,388
EMEA	HR	4,316,700	3,679,388
EMEA	Marketing	2,274,878	3,679,388

Arguments

measure

The measure that you want to do the calculation for, for example `sum( {Sales Amt} )`. Use an aggregation if the calculation level is set to `NULL` or `POST_AGG_FILTER`. Don't use an aggregation if the calculation level is set to `PRE_FILTER` or `PRE_AGG`.

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in `{ }` (curly braces), if it is more than one word. The entire list is enclosed in `[ ]` (square brackets).

calculation level

(Optional) Specifies the calculation level to use:

- **PRE_FILTER** – Prefilter calculations are computed before the dataset filters.
- **PRE_AGG** – Preaggregate calculations are computed before applying aggregations and top and bottom *N* filters to the visuals.
- **POST_AGG_FILTER** – (Default) Table calculations are computed when the visuals display.

This value defaults to `POST_AGG_FILTER` when blank. For more information, see [Using level-aware calculations in Quick](#).

Example

The following example gets the average `sum(Sales)` partitioned over `City` and `State`.

```
avgOver
(
    sum(Sales),
    [City, State]
)
```

countOver

The `countOver` function calculates the count of a dimension or measure partitioned by a list of dimensions.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
countOver  
(  
  measure or dimension field  
  , [ partition_field, ... ]  
  , calculation level  
)
```

Arguments

measure or dimension field

The measure or dimension that you want to do the calculation for, for example `sum( {Sales Amt} )`. Use an aggregation if the calculation level is set to `NULL` or `POST_AGG_FILTER`. Don't use an aggregation if the calculation level is set to `PRE_FILTER` or `PRE_AGG`.

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in `{ }` (curly braces), if it is more than one word. The entire list is enclosed in `[ ]` (square brackets).

calculation level

(Optional) Specifies the calculation level to use:

- **PRE_FILTER** – Prefilter calculations are computed before the dataset filters.
- **PRE_AGG** – Preaggregate calculations are computed before applying aggregations and top and bottom *N* filters to the visuals.
- **POST_AGG_FILTER** – (Default) Table calculations are computed when the visuals display.

This value defaults to `POST_AGG_FILTER` when blank. For more information, see [Using level-aware calculations in Quick](#).

Example

The following example gets the count of Sales partitioned over City and State.

```
countOver
(
  Sales,
  [City, State]
)
```

The following example gets the count of {County} partitioned over City and State.


```
countOver
(
  {County},
  [City, State]
)
```

The following example shows the count of Billed Amount over Customer Region. The fields in the table calculation are in the field wells of the visual.

```
countOver
(
  sum({Billed Amount}),
  [{Customer Region}]
)
```

The following screenshot shows the results of the example. Because there are no other fields involved, the count is one for each region.

Field wells



If you add additional fields, the count changes. In the following screenshot, we add Customer Segment and Service Line. Each of those fields contains three unique values. With 3 segments, 3 service lines, and 3 regions, the calculated field shows 9.

Field wells

The screenshot shows two panels. The 'Group by' panel on the left has a light gray header and three blue buttons with white text and downward arrows: 'Customer Segment', 'Service Line', and 'Customer Region'. The 'Value' panel on the right has a light gray header and a single green button with white text and a downward arrow: 'countOver (Custom)'.

Countover by Customer Segment, Service Line, and Customer Region			
Customer Segment	Service Line	Customer Region	countOver
Enterprise	Billing	APAC	9
Enterprise	Billing	EMEA	9
Enterprise	Billing	US	9
Enterprise	HR	APAC	9

If you add the two additional fields to the partitioning fields in the calculated field, `countOver( sum({Billed Amount}), [{Customer Region}, {Customer Segment}, {Service Line}])`, then the count is again 1 for each row.

Field wells

Group by

Customer Segment

Service Line

Customer Region

Value

countOver (Custom)

Countover by Customer Region, Service Line, and Customer Segment

Customer Segment	Service Line	Customer Region	countOver
Enterprise	Billing	APAC	1
Enterprise	Billing	EMEA	1
Enterprise	Billing	US	1
Enterprise	HR	APAC	1
Enterprise	HR	EMEA	1
Enterprise	HR	US	1

maxOver

The `maxOver` function calculates the maximum of a measure or date partitioned by a list of dimensions.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
maxOver
(
    measure
    , [ partition_field, ... ]
    , calculation level
)
```

Arguments

measure

The measure that you want to do the calculation for, for example `sum( {Sales Amt} )`. Use an aggregation if the calculation level is set to `NULL` or `POST_AGG_FILTER`. Don't use an aggregation if the calculation level is set to `PRE_FILTER` or `PRE_AGG`.

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in `{ }` (curly braces), if it is more than one word. The entire list is enclosed in `[ ]` (square brackets).

calculation level

(Optional) Specifies the calculation level to use:

- **PRE_FILTER** – Prefilter calculations are computed before the dataset filters.
- **PRE_AGG** – Preaggregate calculations are computed before applying aggregations and top and bottom *N* filters to the visuals.
- **POST_AGG_FILTER** – (Default) Table calculations are computed when the visuals display.

This value defaults to `POST_AGG_FILTER` when blank. For more information, see [Using level-aware calculations in Quick](#).

Example

The following example calculates the maximum `sum(Sales)`, partitioned by `City` and `State`.

```
maxOver
(  
    sum(Sales),  
    [City, State]  
)
```

The following example shows the maximum `Billed Amount` over `Customer Region`. The fields in the table calculation are in the field wells of the visual.

```
maxOver
```

```
(
    sum({Billed Amount}),
    [{Customer Region}]
)
```

The following screenshot shows the results of the example. With the addition of Service Line, the total amount billed for each is displayed, and the maximum of these three values displays in the calculated field.

Field wells

Group by

Customer Region

Service Line

Value

Billed Amount (Sum)

maxOver (Custom)

Sum of Billed Amount and Maxover by Customer Region and Service Line

Customer Region	Service Line	Billed Amount	maxOver
APAC	Billing	3,569,780	3,569,780
APAC	HR	3,441,106	= 3,569,780
APAC	Marketing	1,379,768	3,569,780
EMEA	Billing	4,446,586	4,446,586
EMEA	HR	4,316,700	4,446,586
EMEA	Marketing	2,274,878	4,446,586

minOver

The minOver function calculates the minimum of a measure or date partitioned by a list of dimensions.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
minOver
(
    measure
)
```

```
,[ partition_field, ... ]  
,calculation level  
)
```

Arguments

measure

The measure that you want to do the calculation for, for example `sum({Sales Amt})`. Use an aggregation if the calculation level is set to `NULL` or `POST_AGG_FILTER`. Don't use an aggregation if the calculation level is set to `PRE_FILTER` or `PRE_AGG`.

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in `{ }` (curly braces), if it is more than one word. The entire list is enclosed in `[ ]` (square brackets).

calculation level

(Optional) Specifies the calculation level to use:

- **PRE_FILTER** – Prefilter calculations are computed before the dataset filters.
- **PRE_AGG** – Preaggregate calculations are computed before applying aggregations and top and bottom *N* filters to the visuals.
- **POST_AGG_FILTER** – (Default) Table calculations are computed when the visuals display.

This value defaults to `POST_AGG_FILTER` when blank. For more information, see [Using level-aware calculations in Quick](#).

Example

The following example calculates the min `sum(Sales)`, partitioned by `City` and `State`.

```
minOver  
(  
    sum(Sales),  
    [City, State]  
)
```

The following example shows the minimum Billed Amount over Customer Region. The fields in the table calculation are in the field wells of the visual.

```
minOver
(
    sum({Billed Amount}),
    [{Customer Region}]
)
```

The following screenshot shows the results of the example. With the addition of Service Line, the total amount billed for each is displayed, and the minimum of these three values displays in the calculated field.

Field wells

Group by

Customer Region

Service Line

Value

Billed Amount (Sum)

minOver (Custom)

Customer Region	Service Line	Billed Amount	minOver
APAC	Billing	3,569,780	1,379,768
APAC	HR	3,441,106	= 1,379,768
APAC	Marketing	1,379,768	1,379,768
EMEA	Billing	4,446,586	2,274,878
EMEA	HR	4,316,700	2,274,878
EMEA	Marketing	2,274,878	2,274,878

percentileOver

The `percentileOver` function calculates the n th percentile of a measure partitioned by a list of dimensions. There are two varieties of the `percentileOver` calculation available in Quick:

- [percentileContOver](#) uses linear interpolation to determine result.
- [percentileDiscOver](#) uses actual values to determine result.

The `percentileOver` function is an alias of `percentileDiscOver`.

percentileContOver

The `percentileContOver` function calculates the percentile based on the actual numbers in measure. It uses the grouping and sorting that are applied in the field wells. The result is partitioned by the specified dimension at the specified calculation level.

Use this function to answer the following question: Which actual data points are present in this percentile? To return the nearest percentile value that is present in your dataset, use `percentileDiscOver`. To return an exact percentile value that might not be present in your dataset, use `percentileContOver` instead.

Syntax

```
percentileContOver (  
    measure  
    , percentile-n  
    , [partition-by, ...]  
    , calculation-level  
)
```

Arguments

measure

Specifies a numeric value to use to compute the percentile. The argument must be a measure or metric. Nulls are ignored in the calculation.

percentile-n

The percentile value can be any numeric constant 0–100. A percentile value of 50 computes the median value of the measure.

partition-by

(Optional) One or more dimensions that you want to partition by, separated by commas. Each field in the list is enclosed in { } (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

calculation-level

Specifies where to perform the calculation in relation to the order of evaluation. There are three supported calculation levels:

- PRE_FILTER
- PRE_AGG
- POST_AGG_FILTER (default) – To use this calculation level, specify an aggregation on measure, for example `sum(measure)`.

PRE_FILTER and PRE_AGG are applied before the aggregation occurs in a visualization. For these two calculation levels, you can't specify an aggregation on measure in the calculated field expression. To learn more about calculation levels and when they apply, see [Order of evaluation in Amazon Quick](#) and [Using level-aware calculations in Quick](#).

Returns

The result of the function is a number.

Example of percentileContOver

The following example helps explain how percentileContOver works.

Example Comparing calculation levels for the median

The following example shows the median for a dimension (category) by using different calculation levels with the percentileContOver function. The percentile is 50. The dataset is filtered by a region field. The code for each calculated field is as follows:

- `example = left( category, 1 )` (A simplified example.)
- `pre_agg = percentileContOver ( {Revenue} , 50 , [ example ] , PRE_AGG)`
- `pre_filter = percentileContOver ( {Revenue} , 50 , [ example ] , PRE_FILTER)`
- `post_agg_filter = percentileContOver ( sum ( {Revenue} ) , 50 , [ example ] , POST_AGG_FILTER )`

example	pre_filter	pre_agg	post_agg_filter
0	106,728	119,667	4,117,579
1	102,898	95,946	2,307,547
2	97,807	93,963	554,570
3	101,043	112,585	2,709,057
4	96,533	99,214	3,598,358

5	106,293	97,296	1,875,648
6	97,118	69,159	1,320,672
7	100,201	90,557	969,807

percentileDiscOver

The `percentileDiscOver` function calculates the percentile based on the actual numbers in measure. It uses the grouping and sorting that are applied in the field wells. The result is partitioned by the specified dimension at the specified calculation level. The `percentileOver` function is an alias of `percentileDiscOver`.

Use this function to answer the following question: Which actual data points are present in this percentile? To return the nearest percentile value that is present in your dataset, use `percentileDiscOver`. To return an exact percentile value that might not be present in your dataset, use `percentileContOver` instead.

Syntax

```
percentileDiscOver (  
    measure  
    , percentile-n  
    , [partition-by, ...]  
    , calculation-level  
)
```

Arguments

measure

Specifies a numeric value to use to compute the percentile. The argument must be a measure or metric. Nulls are ignored in the calculation.

percentile-n

The percentile value can be any numeric constant 0–100. A percentile value of 50 computes the median value of the measure.

partition-by

(Optional) One or more dimensions that you want to partition by, separated by commas. Each field in the list is enclosed in { } (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

calculation-level

Specifies where to perform the calculation in relation to the order of evaluation. There are three supported calculation levels:

- PRE_FILTER
- PRE_AGG
- POST_AGG_FILTER (default) – To use this calculation level, you need to specify an aggregation on measure, for example `sum(measure)`.

PRE_FILTER and PRE_AGG are applied before the aggregation occurs in a visualization. For these two calculation levels, you can't specify an aggregation on measure in the calculated field expression. To learn more about calculation levels and when they apply, see [Order of evaluation in Amazon Quick](#) and [Using level-aware calculations in Quick](#).

Returns

The result of the function is a number.

Example of percentileDiscOver

The following example helps explain how percentileDiscOver works.

Example Comparing calculation levels for the median

The following example shows the median for a dimension (category) by using different calculation levels with the percentileDiscOver function. The percentile is 50. The dataset is filtered by a region field. The code for each calculated field is as follows:

- `example = left( category, 1 )` (A simplified example.)
- `pre_agg = percentileDiscOver ( {Revenue} , 50 , [ example ] , PRE_AGG)`
- `pre_filter = percentileDiscOver ( {Revenue} , 50 , [ example ] , PRE_FILTER)`
- `post_agg_filter = percentileDiscOver ( sum ( {Revenue} ) , 50 , [ example ] , POST_AGG_FILTER )`

example	pre_filter	pre_agg	post_agg_filter

0	106,728	119,667	4,117,579
1	102,898	95,946	2,307,547
2	97,629	92,046	554,570
3	100,867	112,585	2,709,057
4	96,416	96,649	3,598,358
5	106,293	97,296	1,875,648
6	97,118	64,395	1,320,672
7	99,915	90,557	969,807

Example The median

The following example calculates the median (the 50th percentile) of Sales partitioned by City and State.

```
percentileDiscOver
(
  Sales,
  50,
  [City, State]
)
```

The following example calculates the 98th percentile of `sum({Billed Amount})` partitioned by Customer Region. The fields in the table calculation are in the field wells of the visual.

```
percentileDiscOver
(
  sum({Billed Amount}),
  98,
  [{Customer Region}]
)
```

The following screenshot shows the how these two examples look on a chart.

Sum of Billed Amount, PercOver50, and PercOver98 by Customer Region			
Customer Region	Billed Amount	percOver50	percOver98
APAC	\$598,114	\$850,406	\$1,648,034
EMEA	\$850,406	\$850,406	\$1,648,034
US	\$1,648,034	\$850,406	\$1,648,034

percentOfTotal

The `percentOfTotal` function calculates the percentage a measure contributes to the total, based on the dimensions specified.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
percentOfTotal  
(  
    measure  
    , [ partition_field, ... ]  
)
```

Arguments

measure

An aggregated measure that you want to see the percent of total for. Currently, the `distinct count` aggregation is not supported for `percentOfTotal`.

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in `{ }` (curly braces), if it is more than one word. The entire list is enclosed in `[ ]` (square brackets).

Example

The following example creates a calculation for the percent of total `Sales` contributed by each `State`.

```
percentOfTotal  
(  
    sum(Sales),  
    [State]  
)
```

The following example calculates the percent that a specific Billed Amount is when compared to the total Billed Amount, partitioned by ([Service Line] ASC). The fields in the table calculation are in the field wells of the visual.

```
percentOfTotal
(
    sum( {Billed Amount} ),
    [{Service Line}]
)
```

The following screenshot shows the results of the example. The red highlights show that the partition field with the value "Billing" has three entries, one for each region. The total billed amount for this service line is divided into three percentages, which total 100 percent. Percentages are rounded and might not always add up to exactly 100 percent.

Field wells

Group by

Service Line

Customer Region

Value

percentOfTotal (Custom)

Billed Amount (Sum)

Service Line	Customer Region	percentOfTotal	Billed Amount
Billing	APAC	20.6%	3,569,779.71
Billing	EMEA	25.6%	4,446,586.13
Billing	US	53.8%	9,330,832.51
HR	APAC	20.0%	3,441,106.16
HR	EMEA	25.1%	4,316,700.48
HR	US	55.0%	9,464,168.33

periodOverPeriodDifference

The periodOverPeriodDifference function calculates the difference of a measure over two different time periods as specified by period granularity and offset. Unlike a difference calculation,

this function uses a date-based offset instead of a fixed sized offset. This ensures that only the correct dates are compared, even if data points are missing in the dataset.

Syntax

```
periodOverPeriodDifference(  
  measure,  
  date,  
  period,  
  offset)
```

Arguments

measure

An aggregated measure that you want to perform the periodOverPeriod calculation on.

dateTime

The Date dimension over which we are computing Period-Over-Period calculations.

period

(Optional) The time period across which you're computing the computation. Granularity of YEAR means YearToDate computation, Quarter means QuarterToDate, and so on. Valid granularities include YEAR, QUARTER, MONTH, WEEK, DAY, HOUR, MINUTE, and SECONDS.

The defaults value is the visual date dimension granularity.

offset

(Optional) The offset can be a positive or negative integer representing the prior time period (specified by period) that you want to compare against. For instance, period of a quarter with offset 1 means comparing against the previous quarter.

The default value is 1.

Example

The following example uses a calculated field PeriodOverPeriod to display the sales amount difference of yesterday

```
periodOverPeriodDifference(sum(Sales), {Order Date})
```

The screenshot shows the Amazon QuickSight interface. In the 'Group by' section, 'Order Date' is selected. In the 'Value' section, 'Sales (Sum)' and 'PeriodOverPeriod (Custom)' are listed. Below the fields, the calculated field 'PeriodOverPeriodDifference' is shown. The resulting table displays the sum of sales and the period-over-period difference for each order date.

Order Date	Sales	PeriodOverPeriod
Jan 4, 2020	16	
Jan 5, 2020	288	272
Jan 6, 2020	20	-269
Jan 7, 2020	4,407	4,388
Jan 8, 2020	87	-4,320

The following example uses a calculated field `PeriodOverPeriod` to display the sales amount difference of previous 2 months. Below example is comparing sales of Mar2020 with Jan2020.

```
periodOverPeriodDifference(sum(Sales), {Order Date}, MONTH, 1)
```

The screenshot shows the Amazon QuickSight interface. In the 'Group by' section, 'Order Date (MONTH)' is selected. In the 'Value' section, 'Sales (Sum)' and 'PeriodOverPeriod2 (Custom)' are listed. Below the fields, the calculated field 'PeriodOverPeriodDifference' is shown. The resulting table displays the sum of sales and the period-over-period difference for each month, comparing March 2020 with January 2020.

Order ...	Sales	PeriodOverPeriod2
Jan 2020	13,946	
Feb 2020	4,811	
Mar 2020	55,691	41,745
Apr 2020	28,295	23,485
May 2020	23,648	-32,043

periodOverPeriodLastValue

The `periodOverPeriodLastValue` function calculates the last (previous) value of a measure from the previous time period as specified by the period granularity and offset. This function uses a

date-based offset instead of a fixed sized offset. This ensures only the correct dates are compared, even if data points are missing in the dataset.

Syntax

```
periodOverPeriodLastValue(  
    measure,  
    date,  
    period,  
    offset)
```

Arguments

measure

An aggregated measure that you want to see the difference for.

date

The date dimension over which you're computing periodOverPeriod calculations.

period

(Optional) The time period across which you're computing the computation. Granularity of YEAR means YearToDate computation, Quarter means QuarterToDate, and so on. Valid granularities include YEAR, QUARTER, MONTH, WEEK, DAY, HOUR, MINUTE, and SECONDS.

This argument defaults to the granularity of the visual aggregation

offset

(Optional) The offset can a positive or negative integer representing the prior time period (specified by period) that you want to compare against. For instance, period of a quarter with offset 1 means comparing against the previous quarter.

This argument default value is 1.


Example

The following example calculates the month over month value in sales with the visual dimension granularity and default offset of 1.

```
periodOverPeriodLastValue(sum(Sales), {Order Date})
```

The following example calculates the month over month value in sales with a fixed granularity of MONTH and fixed offset of 1.

```
periodOverPeriodLastValue(sum(Sales), {Order Date},MONTH, 1)
```



periodOverPeriodPercentDifference

The `periodOverPeriodPercentDifference` function calculates the percent difference of a measure over two different time periods as specified by the period granularity and offset. Unlike `percentDifference`, this function uses a date-based offset instead of a fixed sized offset. This ensures only the correct dates are compared, even if data points are missing in the dataset.

Syntax

```
periodOverPeriodPercentDifference(
```



```
measure,  
date,  
period,  
offset)
```

Arguments

measure

An aggregated measure that you want to see the difference for.

date

The date dimension over which you're computing periodOverPeriod calculations.

period

(Optional) The time period across which you're computing the computation. Granularity of YEAR means YearToDate computation, Quarter means QuarterToDate, and so on. Valid granularities include YEAR, QUARTER, MONTH, WEEK, DAY, HOUR, MINUTE, and SECONDS.

This argument defaults to the granularity of the visual aggregation

offset

(Optional) The offset can a positive or negative integer representing the prior time period (specified by period) that you want to compare against. For instance, period of a quarter with offset 1 means comparing against the previous quarter.

This argument default value is 1.

Example

The following example calculates the month over month percent difference in sales with the visual dimension granularity and default offset of 1.

```
periodOverPeriodPercentDifference(sum(Sales), {Order Date})
```

The following example calculates the month over month percent difference in sales with a fixed granularity of MONTH and fixed offset of 1.

```
periodOverPeriodPercentDifference(sum(Sales), {Order Date}, MONTH, 1)
```

Field wells
edit, replace, and remove datasets.

Group by

Order Date (MONTH) ▾

Value

MoMPercentDifference (Custom) ▾

Sales (Sum) ▾

PeriodOverPeriod | PTDOverTime / PTD Agg | **PeriodOverPeriod** ▾ +

Sum of Sales and Mompercentdifference by Order Date

Order Date	MoMPercentDifference	Sales
Jan 2015		274,766.92
Feb 2015	18.68%	326,101.47
Mar 2015	-16.68%	271,696.67
Apr 2015	43.48%	389,831.95
May 2015	-21.36%	306,572.07
Jun 2015	15.92%	355,368.8

periodToDateAvgOverTime

The `periodToDateAvgOverTime` function calculates the average of a measure for a given time granularity (for instance, a quarter) up to a point in time.

Syntax

```
periodToDateAvgOverTime(
  measure,
  dateTime,
  period)
```

Arguments

measure

An aggregated measure that you want to do the calculation

dateTime

The date dimension over which you're computing PeriodOverTime calculations.

period

(Optional) The time period across which you're computing the computation. Granularity of YEAR means YearToDate computation, Quarter means QuarterToDate, and so on. Valid granularities include YEAR, QUARTER, MONTH, WEEK, DAY, HOUR, MINUTE, and SECONDS.

The default value is the visual's date dimension granularity.

Example

The following function calculates the average fare amount month over month.

```
periodToDateAvgOverTime(sum({fare_amount}), pickupDatetime, MONTH)
```

Field wells

Group by

pickupDatetime (WEEK) ▾

Value

fare_amount (Sum) ▾

PTDOverTimeAvg (Custom) ▾

Sheet 1 ▾ +

Ptdovertimeavg and Sum of Fare_amount by Pickupdatetime

pickupDatetime	fare_amount	PTDOverTimeAvg
Jun 27, 2021	D 20276346.25	31143423.01
Jun 20, 2021	C 34294039.05	34765781.93
Jun 13, 2021	B 34311811.14	35001653.37
Jun 6, 2021	A 35691495.60	35691495.60
May 30, 2021	33927942.61	35770454.41
May 23, 2021	32781949.16	36231082.36
May 16, 2021	38225816.10	37380793.42

$(A+B+C+D)/4$

$(A+B+C)/3$

$(A+B)/2$

$A/1$

periodToDateCountOverTime

The `periodToDateCountOverTime` function calculates the count of a dimension or measure for a given time granularity (for instance, a quarter) up to a point in time.

Syntax

```
periodToDateCountOverTime(
  measure,
  dateTime,
  period)
```

Arguments

measure

An aggregated measure that you want to do the calculation

dateTime

The date dimension over which you're computing PeriodOverTime calculations.

period

(Optional) The time period across which you're computing the computation. Granularity of YEAR means YearToDate computation, Quarter means QuarterToDate, and so on. Valid granularities include YEAR, QUARTER, MONTH, WEEK, DAY, HOUR, MINUTE, and SECONDS.

The default value is the visual's date dimension granularity.

Example

The following example calculates the count of vendors month over month.

```
periodToDateCountOverTime(count(vendorid), pickupDatetime, MONTH)
```

Field wells

Group by

pickupDatetime (WEEK)

vendorid

Value

PTDOverTimeCount (Custom)

Sheet 1

+

Ptdovertimecount by Pickupdatetime and Vendorid

pickupDatetime	vendorid	PTDOverTimeCount
Jun 27, 2021	1	4
Jun 27, 2021	2	4
Jun 20, 2021	1	3
Jun 20, 2021	2	3
Jun 13, 2021	1	2
Jun 13, 2021	2	2

periodToDateMaxOverTime

The `periodToDateMaxOverTime` function calculates the maximum of a measure for a given time granularity (for instance, a quarter) up to a point in time.

Syntax

```
periodToDateMaxOverTime(  
  measure,  
  dateTime,  
  period)
```

Arguments

measure

An aggregated measure that you want to do the calculation

dateTime

The date dimension over which you're computing PeriodOverTime calculations.

period

(Optional) The time period across which you're computing the computation. Granularity of YEAR means YearToDate computation, Quarter means QuarterToDate, and so on. Valid granularities include YEAR, QUARTER, MONTH, WEEK, DAY, HOUR, MINUTE, and SECONDS.

The default value is the visual's date dimension granularity.

Example

The following example calculates the maximum fare amount month over month.

```
periodToDateMaxOverTime(max({fare_amount}), pickupDatetime, MONTH)
```

Field wells

Group by

pickupDatetime (WEEK) ▾

Value

fare_amount (Max) ▾

PTDOverTimeMax (Custom) ▾

Sheet 1 ▾ +

Max of Fare_amount and PtdovertimeMax by Pickupdatetime

pickupDatetime	fare_amount	PTDOverTimeMax
Jun 27, 2021	628544.74	628,544.74
Jun 20, 2021	8007.00	187,440.96
Jun 13, 2021	8452.00	187,440.96
Jun 6, 2021	187440.96	187,440.96
May 30, 2021	133057.84	133,057.84
May 23, 2021	4886.00	8,007
May 16, 2021	1520.40	8,007

periodToDateMinOverTime

The `periodToDateMinOverTime` function calculates the minimum of a measure for a given time granularity (for instance, a quarter) up to a point in time.

Syntax

```
periodToDateMinOverTime(
  measure,
  dateTime,
  period)
```

Arguments

measure

An aggregated measure that you want to do the calculation

dateTime

The date dimension over which you're computing PeriodOverTime calculations.

period

(Optional) The time period across which you're computing the computation. Granularity of YEAR means YearToDate computation, Quarter means QuarterToDate, and so on. Valid granularities include YEAR, QUARTER, MONTH, WEEK, DAY, HOUR, MINUTE, and SECONDS.

The default value is the visual's date dimension granularity.

Example

The following example calculates the minimum fare amount month over month.

```
periodToDateMinOverTime(min({fare_amount}), pickupDatetime, MONTH)
```

Field wells

Group by

pickupDatetime (WEEK) ▾

Value

fare_amount (Min) ▾

PTDOverTimeMin (Custom) ▾

Sheet 1 ▾ +

Min of Fare_amount and Ptdovertimemin by Pickupdatetime

pickupDatetime	fare_amount	PTDOverTimeMin
Jun 27, 2021	-250.00	-450
Jun 20, 2021	-450.00	-450
Jun 13, 2021	-273.00	-287
Jun 6, 2021	-287.00	-287
May 30, 2021	-199.00	-410
May 23, 2021	-300.00	-410

periodToDateSumOverTime

The `periodToDateSumOverTime` function calculates the sum of a measure for a given time granularity (for instance, a quarter) up to a point in time.

Syntax

```
periodToDateSumOverTime(  
    measure,  
    dateTime,  
    period)
```

Arguments

measure

An aggregated measure that you want to do the calculation

dateTime

The date dimension over which you're computing PeriodOverTime calculations.

period

(Optional) The time period across which you're computing the computation. Granularity of YEAR means YearToDate computation, Quarter means QuarterToDate, and so on. Valid granularities include YEAR, QUARTER, MONTH, WEEK, DAY, HOUR, MINUTE, and SECONDS.

The default value is the visual's date dimension granularity.

Example

The following function returns the total fare amount month over month.

```
periodToDateSumOverTime(sum({fare_amount}), pickupDatetime, MONTH)
```

Field wells

Group by

pickupDatetime (WEEK)

Value

fare_amount (Sum)

PTDOverTimeSum (Custom)

Sheet 1

+

Sum of Fare_amount and Ptdovertimesum by Pickupdatetime

pickupDatetime		fare_amount	PTDOverTimeSum	
Jun 27, 2021	D	20,276,346.25	124,573,692.04	A+B+C+D
Jun 20, 2021	C	34,294,039.05	104,297,345.79	A+B+C
Jun 13, 2021	B	34,311,811.14	70,003,306.74	A+B
Jun 6, 2021	A	35,691,495.6	35,691,495.6	A
May 30, 2021		33,927,942.61	178,852,272.03	
May 23, 2021		32,781,949.16	144,924,329.42	
May 16, 2021		38,225,816.1	112,142,380.26	
May 9, 2021		36,938,239.42	73,916,564.16	

stdevOver

The stdevOver function calculates the standard deviation of the specified measure, partitioned by the chosen attribute or attributes, based on a sample.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
stdevOver
(
    measure
    , [ partition_field, ... ]
    , calculation level
)
```

Arguments

measure

The measure that you want to do the calculation for, for example `sum({Sales Amt})`. Use an aggregation if the calculation level is set to `NULL` or `POST_AGG_FILTER`. Don't use an aggregation if the calculation level is set to `PRE_FILTER` or `PRE_AGG`.

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in `{ }` (curly braces), if it is more than one word. The entire list is enclosed in `[ ]` (square brackets).

calculation level

(Optional) Specifies the calculation level to use:

- **PRE_FILTER** – Prefilter calculations are computed before the dataset filters.
- **PRE_AGG** – Preaggregate calculations are computed before applying aggregations and top and bottom *N* filters to the visuals.
- **POST_AGG_FILTER** – (default) table calculations are computed when the visuals display.

This value defaults to `POST_AGG_FILTER` when blank. For more information, see [Using level-aware calculations in Quick](#).

Example

The following example calculates the standard deviation of `sum(Sales)`, partitioned by `City` and `State`, based on a sample..

```
stdevOver
(  
    sum(Sales),  
    [City, State]  
)
```

The following example calculates the standard deviation of `Billed Amount` over `Customer Region`, based on a sample. The fields in the table calculation are in the field wells of the visual.

```
stdevOver
```

```
(  
    sum({Billed Amount}),  
    [{Customer Region}]  
)
```

stdevpOver

The `stdevpOver` function calculates the standard deviation of the specified measure, partitioned by the chosen attribute or attributes, based on a biased population.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
stdevpOver  
(  
    measure  
    , [ partition_field, ... ]  
    , calculation level  
)
```

Arguments

measure

The measure that you want to do the calculation for, for example `sum({Sales Amt})`. Use an aggregation if the calculation level is set to `NULL` or `POST_AGG_FILTER`. Don't use an aggregation if the calculation level is set to `PRE_FILTER` or `PRE_AGG`.

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in `{ }` (curly braces), if it is more than one word. The entire list is enclosed in `[ ]` (square brackets).

calculation level

(Optional) Specifies the calculation level to use:

- **PRE_FILTER** – Prefilter calculations are computed before the dataset filters.
- **PRE_AGG** – Preaggregate calculations are computed before applying aggregations and top and bottom *N* filters to the visuals.

- **POST_AGG_FILTER** – (default) table calculations are computed when the visuals display.

This value defaults to `POST_AGG_FILTER` when blank. For more information, see [Using level-aware calculations in Quick](#).

Example

The following example calculates the standard deviation of `sum(Sales)`, partitioned by `City` and `State`, based on a biased population.

```
stdevpOver
(  
    sum(Sales),  
    [City, State]  
)
```

The following example calculates the standard deviation of `Billed Amount` over `Customer Region`, based on a biased population. The fields in the table calculation are in the field wells of the visual.

```
stdevpOver
(  
    sum({Billed Amount}),  
    [{Customer Region}]  
)
```

varOver

The `varOver` function calculates the variance of the specified measure, partitioned by the chosen attribute or attributes, based on a sample.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
varOver
(  
    measure  
    , [ partition_field, ... ]  
    , calculation level
```

```
)
```

Arguments

measure

The measure that you want to do the calculation for, for example `sum({Sales Amt})`. Use an aggregation if the calculation level is set to `NULL` or `POST_AGG_FILTER`. Don't use an aggregation if the calculation level is set to `PRE_FILTER` or `PRE_AGG`.

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in `{ }` (curly braces), if it is more than one word. The entire list is enclosed in `[ ]` (square brackets).

calculation level

(Optional) Specifies the calculation level to use:

- **PRE_FILTER** – Prefilter calculations are computed before the dataset filters.
- **PRE_AGG** – Preaggregate calculations are computed before applying aggregations and top and bottom *N* filters to the visuals.
- **POST_AGG_FILTER** – (Default) Table calculations are computed when the visuals display.

This value defaults to `POST_AGG_FILTER` when blank. For more information, see [Using level-aware calculations in Quick](#).

Example

The following example calculates the variance of `sum(Sales)`, partitioned by `City` and `State`, based on a sample.

```
varOver
(  
    sum(Sales),  
    [City, State]  
)
```

The following example calculates the variance of `Billed Amount` over `Customer Region`, based on a sample. The fields in the table calculation are in the field wells of the visual.

```
varOver
(
    sum({Billed Amount}),
    [{Customer Region}]
)
```

varpOver

The `varpOver` function calculates the variance of the specified measure, partitioned by the chosen attribute or attributes, based on a biased population.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
varpOver
(
    measure
    , [ partition_field, ... ]
    , calculation level
)
```

Arguments

measure

The measure that you want to do the calculation for, for example `sum({Sales Amt})`. Use an aggregation if the calculation level is set to `NULL` or `POST_AGG_FILTER`. Don't use an aggregation if the calculation level is set to `PRE_FILTER` or `PRE_AGG`.

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in `{ }` (curly braces), if it is more than one word. The entire list is enclosed in `[ ]` (square brackets).

calculation level

(Optional) Specifies the calculation level to use:

- **PRE_FILTER** – Prefilter calculations are computed before the dataset filters.

- **PRE_AGG** – Preaggregate calculations are computed before applying aggregations and top and bottom *N* filters to the visuals.
- **POST_AGG_FILTER** – (Default) Table calculations are computed when the visuals display.

This value defaults to POST_AGG_FILTER when blank. For more information, see [Using level-aware calculations in Quick](#).

Example

The following example calculates the variance of `sum(Sales)`, partitioned by `City` and `State`, based on a biased population.

```
varpOver  
(  
    sum(Sales),  
    [City, State]  
)
```

The following example calculates the variance of `Billed Amount` over `Customer Region`, based on a biased population. The fields in the table calculation are in the field wells of the visual.

```
varpOver  
(  
    sum({Billed Amount}),  
    [{Customer Region}]  
)
```

sumOver

The `sumOver` function calculates the sum of a measure partitioned by a list of dimensions.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
sumOver  
(  
    measure  
    , [ partition_field, ... ]  
    , calculation level  
)
```


Arguments

measure

The measure that you want to do the calculation for, for example `sum( {Sales Amt} )`. Use an aggregation if the calculation level is set to `NULL` or `POST_AGG_FILTER`. Don't use an aggregation if the calculation level is set to `PRE_FILTER` or `PRE_AGG`.

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in `{ }` (curly braces), if it is more than one word. The entire list is enclosed in `[ ]` (square brackets).

calculation level

(Optional) Specifies the calculation level to use:

- **PRE_FILTER** – Prefilter calculations are computed before the dataset filters.
- **PRE_AGG** – Preaggregate calculations are computed before applying aggregations and top and bottom *N* filters to the visuals.
- **POST_AGG_FILTER** – (default) table calculations are computed when the visuals display.

This value defaults to `POST_AGG_FILTER` when blank. For more information, see [Using level-aware calculations in Quick](#).

Example

The following example calculates the sum of `sum(Sales)`, partitioned by `City` and `State`.

```
sumOver
(  
    sum(Sales),  
    [City, State]  
)
```

The following example sums `Billed Amount` over `Customer Region`. The fields in the table calculation are in the field wells of the visual.

```
sumOver
```

```
(  
    sum({Billed Amount}),  
    [{Customer Region}]  
)
```

The following screenshot shows the results of the example. With the addition of Customer Segment, the total amount billed for each is summed for the Customer Region, and displays in the calculated field.

Field wells

Group by

Customer Region

Customer Segment

Value

Billed Amount (Sum)

sumOver (Custom)

Sum of Billed Amount and Sumover by Customer Region and Customer Segment

Customer Region	Customer Segment	Billed Amount	sumOver
APAC	Enterprise	2,035,949	8,390,654
APAC	SMB	660,097	= 8,390,654
APAC	Startup	5,694,609	8,390,654
EMEA	Enterprise	5,678,783	11,038,164
EMEA	SMB	1,341,834	11,038,164
EMEA	Startup	4,017,547	11,038,164

denseRank

The denseRank function calculates the rank of a measure or a dimension in comparison to the specified partitions. It counts each item only once, ignoring duplicates, and assigns a rank "without holes" so that duplicate values share the same rank.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
denseRank  
(  
    [ sort_order_field ASC_or_DESC, ... ]
```

```
, [ partition_field, ... ]  
)
```

Arguments

sort order field

One or more aggregated fields, either measures or dimensions or both, that you want to sort the data by, separated by commas. You can either specify ascending (**ASC**) or descending (**DESC**) sort order.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

calculation level

(Optional) Specifies the calculation level to use:

- **PRE_FILTER** – Prefilter calculations are computed before the dataset filters.
- **PRE_AGG** – Preaggregate calculations are computed before applying aggregations and top and bottom *N* filters to the visuals.
- **POST_AGG_FILTER** – (Default) Table calculations are computed when the visuals display.

This value defaults to POST_AGG_FILTER when blank. For more information, see [Using level-aware calculations in Quick](#).

Example

The following example densely ranks max(Sales), based on a descending sort order, by State and City. Any cities with the same max(Sales) are assigned the same rank, and the next city is ranked consecutively after them. For example, if three cities share the same ranking, the fourth city is ranked as second.

```
denseRank
```

```
(  
  [max(Sales) DESC],  
  [State, City]  
)
```

The following example densely ranks `max(Sales)`, based on a descending sort order, by `State`. Any states with the same `max(Sales)` are assigned the same rank, and the next is ranked consecutively after them. For example, if three states share the same ranking, the fourth state is ranked as second.

```
denseRank  
(  
  [max(Sales) DESC],  
  [State]  
)
```

rank

The `rank` function calculates the rank of a measure or a dimension in comparison to the specified partitions. It counts each item, even duplicates, once and assigns a rank "with holes" to make up for duplicate values.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
rank  
(  
  [ sort_order_field ASC_or_DESC, ... ]  
  , [ partition_field, ... ]  
)
```

Arguments

sort_order_field

One or more aggregated measures and dimensions that you want to sort the data by, separated by commas. You can specify either ascending (**ASC**) or descending (**DESC**) sort order.

Each field in the list is enclosed in `{ }` (curly braces), if it is more than one word. The entire list is enclosed in `[ ]` (square brackets).

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

calculation level

(Optional) Specifies the calculation level to use:

- **PRE_FILTER** – Prefilter calculations are computed before the dataset filters.
- **PRE_AGG** – Preaggregate calculations are computed before applying aggregations and top and bottom *N* filters to the visuals.
- **POST_AGG_FILTER** – (Default) Table calculations are computed when the visuals display.

This value defaults to POST_AGG_FILTER when blank. For more information, see [Using level-aware calculations in Quick](#).

Example

The following example ranks `max(Sales)`, based on a descending sort order, by State and City, within the State of **WA**. Any cities with the same `max(Sales)` are assigned the same rank, but the next rank includes the count of all previously existing ranks. For example, if three cities share the same ranking, the fourth city is ranked as fourth.

```
rank
(  
  [max(Sales) DESC],  
  [State, City]  
)
```

The following example ranks `max(Sales)`, based on an ascending sort order, by State. Any states with the same `max(Sales)` are assigned the same rank, but the next rank includes the count of all previously existing ranks. For example, if three states share the same ranking, the fourth state is ranked as fourth.

```
rank
(  
  [max(Sales) ASC],  
  [State]
```

```
)
```

The following example ranks Customer Region by total Billed Amount. The fields in the table calculation are in the field wells of the visual.

```
rank(  
  [sum({Billed Amount}) DESC]  
)
```

The following screenshot shows the results of the example, along with the total Billed Amount so you can see how each region ranks.

Field wells

Group by

Customer Region

Value

rank (Custom)

Billed Amount (Sum)

Rank and Sum of Billed Amount by Customer Region

Customer Region	rank	Billed Amount
APAC	3	8,390,654
EMEA	2	11,038,164
US	1	23,547,165

percentileRank

The percentileRank function calculates the percentile rank of a measure or a dimension in comparison to the specified partitions. The percentile rank value(*x*) indicates that the current item is above *x*% of values in the specified partition. The percentile rank value ranges from 0 (inclusive) to 100 (exclusive).

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
percentileRank
```

```
(  
    [ sort_order_field ASC_or_DESC, ... ]  
    , [ {partition_field}, ... ]  
)
```

Arguments

sort order field

One or more aggregated measures and dimensions that you want to sort the data by, separated by commas. You can specify either ascending (**ASC**) or descending (**DESC**) sort order.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

calculation level

(Optional) Specifies the calculation level to use:

- **PRE_FILTER** – Prefilter calculations are computed before the dataset filters.
- **PRE_AGG** – Preaggregate calculations are computed before applying aggregations and top and bottom *N* filters to the visuals.
- **POST_AGG_FILTER** – (Default) Table calculations are computed when the visuals display.

This value defaults to POST_AGG_FILTER when blank. For more information, see [Using level-aware calculations in Quick](#).

Example


The following example does a percentile ranking of max(Sales) in descending order, by State.

```
percentileRank  
(  
    [max(Sales) DESC],  
    [State]  
)
```

The following example does a percentile ranking of Customer Region by total Billed Amount. The fields in the table calculation are in the field wells of the visual.

```
percentileRank(
  [sum({Billed Amount}) DESC],
  [{Customer Region}]
)
```

The following screenshot shows the results of the example, along with the total Billed Amount so you can see how each region compares.



runningAvg

The `runningAvg` function calculates a running average for a measure based on the specified dimensions and sort orders.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
runningAvg
(
  measure
  , [ sortorder_field ASC_or_DESC, ... ]
  , [ partition_field, ... ]
```


)

Arguments

measure

An aggregated measure that you want to see the running average for.

sort order field

One or more measures and dimensions that you want to sort the data by, separated by commas. You can specify either ascending (**ASC**) or descending (**DESC**) sort order.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

Example

The following example calculates a running average of `sum(Sales)`, sorted by Sales, partitioned by City and State.

```
runningAvg
(
  sum(Sales),
  [Sales ASC],
  [City, State]
)
```

The following example calculates a running average of `Billed Amount`, sorted by month (`[truncDate("MM",Date) ASC]`). The fields in the table calculation are in the field wells of the visual.

```
runningAvg
(
  sum({Billed Amount}),
  [truncDate("MM",Date) ASC]
```

```
)
```

runningCount

The `runningCount` function calculates a running count for a measure or dimension, based on the specified dimensions and sort orders.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
runningCount
(  
  measure_or_dimension  
  , [ sortorder_field ASC_or_DESC, ... ]  
  , [ partition_field, ... ]  
)
```

Arguments

measure or dimension

An aggregated measure or dimension that you want to see the running count for.

sort order field

One or more measures and dimensions that you want to sort the data by, separated by commas. You can specify either ascending (**ASC**) or descending (**DESC**) sort order.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

Example

The following example calculates a running count of `sum(Sales)`, sorted by `Sales`, partitioned by `City` and `State`.

```
runningCount
(  
    sum(Sales),  
    [Sales ASC],  
    [City, State]  
)
```

The following example calculates a running count of Billed Amount, sorted by month ([truncDate("MM",Date) ASC]). The fields in the table calculation are in the field wells of the visual.

```
runningCount
(  
    sum({Billed Amount}),  
    [truncDate("MM",Date) ASC]  
)
```

runningMax

The runningMax function calculates a running maximum for a measure based on the specified dimensions and sort orders.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
runningMax
(  
    measure  
    ,[ sortorder_field ASC_or_DESC, ... ]  
    ,[ partition_field, ... ]  
)
```

Arguments

measure

An aggregated measure that you want to see the running maximum for.

sort order field

One or more measures and dimensions that you want to sort the data by, separated by commas. You can specify either ascending (**ASC**) or descending (**DESC**) sort order.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

Example

The following example calculates a running maximum of `sum(Sales)`, sorted by `Sales`, partitioned by `City` and `State`.

```
runningMax
(  
  sum(Sales),  
  [Sales ASC],  
  [City, State]  
)
```

The following example calculates a running maximum of `Billed Amount`, sorted by month (`[truncDate("MM",Date) ASC]`). The fields in the table calculation are in the field wells of the visual.

```
runningMax
(  
  sum({Billed Amount}),  
  [truncDate("MM",Date) ASC]  
)
```

runningMin

The `runningMin` function calculates a running minimum for a measure based on the specified dimensions and sort orders.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
runningMin
(  
  measure  
  , [ sortorder_field ASC_or_DESC, ... ]  
  , [ partition_field, ... ]  
)
```

Arguments

measure

An aggregated measure that you want to see the running minimum for.

sort order field

One or more measures and dimensions that you want to sort the data by, separated by commas. You can specify either ascending (**ASC**) or descending (**DESC**) sort order.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

Example

The following example calculates a running minimum of `sum(Sales)`, sorted by `Sales`, partitioned by `City` and `State`.

```
runningMin
(  
  sum(Sales),  
  [Sales ASC],  
  [City, State]
```

```
)
```

The following example calculates a running minimum of Billed Amount, sorted by month ([truncDate("MM",Date) ASC]). The fields in the table calculation are in the field wells of the visual.

```
runningMin
(  
  sum({Billed Amount}),  
  [truncDate("MM",Date) ASC]  
)
```

runningSum

The runningSum function calculates a running sum for a measure based on the specified dimensions and sort orders.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
runningSum
(  
  measure  
  ,[ sortorder_field ASC_or_DESC, ... ]  
  ,[ partition_field, ... ]  
)
```

Arguments

measure

An aggregated measure that you want to see the running sum for.

sort order field

One or more measures and dimensions that you want to sort the data by, separated by commas. You can specify either ascending (**ASC**) or descending (**DESC**) sort order.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

Example

The following example calculates a running sum of `sum(Sales)`, sorted by Sales, partitioned by City and State.

```
runningSum
(  
  sum(Sales),  
  [Sales ASC],  
  [City, State]  
)
```

The following example calculates a running sum of Billed Amount, sorted by month ([`truncDate("MM",Date) ASC`]). The fields in the table calculation are in the field wells of the visual.

```
runningSum
(  
  sum({Billed Amount}),  
  [truncDate("MM",Date) ASC]  
)
```

The following screenshot shows the results of the example. The red labels show how each amount is added ($a + b = c$) to the next amount, resulting in a new total.

Field wells Rows Date Columns Values Billed Amount (Sum) runningSum (Custom)

Sum of Billed Amount and Runningsum by Date			
Date		Billed Amount	runningSum
Jan 2012		54,675.45	a 54,675.45
Feb 2012	+b	57,127.93	=c 111,803.38
Mar 2012		66,303.97	178,107.35
Apr 2012		66,694.23	244,801.58
May 2012		75,906.62	a 320,708.2
Jun 2012	+b	83,531.67	=c 404,239.87

firstValue

The `firstValue` function calculates the first value of the aggregated measure or dimension partitioned and sorted by specified attributes.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
firstValue
(
    aggregated measure or dimension,
    [ sort_attribute ASC_or_DESC, ... ],
    [ partition_by_attribute, ... ]
)
```

Arguments

aggregated measure or dimension

An aggregated measure or dimension that you want to see the first value for.

sort attribute

One or more aggregated fields, either measures or dimensions or both, that you want to sort the data by, separated by commas. You can either specify ascending (**ASC**) or descending (**DESC**) sort order.

Each field in the list is enclosed in {} (curly braces), if it's more than one word. The entire list is enclosed in [] (square brackets).

partition by attribute

(Optional) One or more measure or dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

Example

The following example calculates the first Destination Airport, sorted by Flight Date, partitioned by Flight Date ascending and Origin Airport.

```
firstValue(  
  {Destination Airport}  
  [{Flight Date} ASC],  
  [  
    {Origin Airport},  
    {Flight Date}  
  ]  
)
```

lastValue

The lastValue function calculates the last value of the aggregated measure or dimension partitioned and sorted by specified attributes.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
lastValue  
(  
  aggregated measure or dimension,  
  [ sort_attribute ASC_or_DESC, ... ],  
  [ partition_by_attribute, ... ]  
)
```

Arguments

aggregated measure or dimension

An aggregated measure or dimension that you want to see the last value for.

sort attribute

One or more aggregated fields, either measures or dimensions or both, that you want to sort the data by, separated by commas. You can either specify ascending (ASC) or descending (DESC) sort order.

Each field in the list is enclosed in {} (curly braces), if it's more than one word. The entire list is enclosed in [] (square brackets).

partition by attribute

(Optional) One or more measures or dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

Example

The following example calculates the last value for `Destination Airport`. This calculation is sorted by the `Flight Date` value and partitioned by the `Flight Date` value sorted in ascending order and the `Origin Airport` value.

```
lastValue(  
  [{Destination Airport}],  
  [{Flight Date} ASC],  
  [  
    {Origin Airport},  
    truncDate('DAY', {Flight Date})  
  ]  
)
```

windowAvg

The `windowAvg` function calculates the average of the aggregated measure in a custom window that is partitioned and sorted by specified attributes. Usually, you use custom window functions

on a time series, where your visual shows a metric and a date field. For example, you can use `windowAvg` to calculate a moving average, which is often used to smooth out the noise in a line chart.

Window functions aren't supported for MySQL versions earlier than 8 and MariaDB versions earlier than 10.2.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
windowAvg
(
    measure
    , [sort_order_field ASC/DESC, ...]
    , start_index
    , end_index
    , [partition_field, ... ]
)
```

Arguments

measure

The aggregated metric that you want to get the average for, for example `sum( {Revenue} )`.

sort attribute

One or more aggregated fields, either measures or dimensions or both, that you want to sort the data by, separated by commas. You can either specify ascending (**ASC**) or descending (**DESC**) sort order.

Each field in the list is enclosed in `{ }` (curly braces), if it's more than one word. The entire list is enclosed in `[ ]` (square brackets).

start index

The start index is a positive integer, indicating *n* rows above the current row. The start index counts the available data points above the current row, rather than counting actual time periods. If your data is sparse (missing months or years, for example), adjust the indexes accordingly.

end index

The end index is a positive integer, indicating n rows below the current row. The end index counts the available data points below the current row, rather than counting actual time periods. If your data is sparse (missing months or years, for example), adjust the indexes accordingly.

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.


Each field in the list is enclosed in {} (curly braces), if it's more than one word. The entire list is enclosed in [] (square brackets).

Example

The following example calculates the moving average of `sum(Revenue)`, partitioned by `SaleDate`. The calculation includes three rows above and two row below of the current row.

```
windowAvg
(
    sum(Revenue),
    [SaleDate ASC],
    3,
    2
)
```

The following screenshot shows the results of this moving average example. The `sum(Revenue)` field is added to the chart to show the difference between the revenue and the moving average of revenue.



windowCount

The `windowCount` function calculates the count of the aggregated measure or dimension in a custom window that is partitioned and sorted by specified attributes. Usually, you use custom window functions on a time series, where your visual shows a metric and a date field.

Window functions aren't supported for MySQL versions earlier than 8 and MariaDB versions earlier than 10.2.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```

windowCount
(
    measure_or_dimension
    , [ sort_order_field ASC/DESC, ... ]
    , start_index
    , end_index
    , [ partition_field, ... ]
  
```

```
)
```

Arguments

measure or dimension

The aggregated metric that you want to get the average for, for example `sum( {Revenue} )`.

sort attribute

One or more aggregated fields, either measures or dimensions or both, that you want to sort the data by, separated by commas. You can either specify ascending (**ASC**) or descending (**DESC**) sort order.

Each field in the list is enclosed in `{ }` (curly braces), if it's more than one word. The entire list is enclosed in `[ ]` (square brackets).

start index

The start index is a positive integer, indicating n rows above the current row. The start index counts the available data points above the current row, rather than counting actual time periods. If your data is sparse (missing months or years, for example), adjust the indexes accordingly.

end index

The end index is a positive integer, indicating n rows below the current row. The end index counts the available data points below the current row, rather than counting actual time periods. If your data is sparse (missing months or years, for example), adjust the indexes accordingly.

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in `{ }` (curly braces), if it's more than one word. The entire list is enclosed in `[ ]` (square brackets).

Example

The following example calculates the moving count of `sum(Revenue)`, partitioned by `SaleDate`. The calculation includes three rows above and two row below of the current row.

```
windowCount
(
    sum(Revenue),
    [SaleDate ASC],
    3,
    2
)
```

windowMax

The windowMax function calculates the maximum of the aggregated measure in a custom window that is partitioned and sorted by specified attributes. Usually, you use custom window functions on a time series, where your visual shows a metric and a date field. You can use windowMax to help you identify the maximum of the metric over a period time.

Window functions aren't supported for MySQL versions earlier than 8 and MariaDB versions earlier than 10.2.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
windowMax
(
    measure
    , [sort_order_field ASC/DESC, ...]
    , start_index
    , end_index
    , [ partition_field, ... ]
)
```

Arguments

measure

The aggregated metric that you want to get the average for, for example sum({Revenue}).

sort attribute

One or more aggregated fields, either measures or dimensions or both, that you want to sort the data by, separated by commas. You can either specify ascending (**ASC**) or descending (**DESC**) sort order.

Each field in the list is enclosed in {} (curly braces), if it's more than one word. The entire list is enclosed in [] (square brackets).

start index

The start index is a positive integer, indicating n rows above the current row. The start index counts the available data points above the current row, rather than counting actual time periods. If your data is sparse (missing months or years, for example), adjust the indexes accordingly.

end index

The end index is a positive integer, indicating n rows below the current row. The end index counts the available data points below the current row, rather than counting actual time periods. If your data is sparse (missing months or years, for example), adjust the indexes accordingly.

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.


Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

Example

The following example calculates the trailing 12-month maximum of `sum(Revenue)`, partitioned by `SaleDate`. The calculation includes 12 rows above and 0 row below of the current row.

```
windowMax
(
    sum(Revenue),
    [SaleDate ASC],
    12,
    0
)
```

The following screenshot shows the results of this trailing 12-month example. The `sum(Revenue)` field is added to the chart to show the difference between the revenue and the trailing 12-month maximum revenue.



windowMin

The `windowMin` function calculates the minimum of the aggregated measure in a custom window that is partitioned and sorted by specified attributes. Usually, you use custom window functions on a time series, where your visual shows a metric and a date field. You can use `windowMin` to help you identify the minimum of the metric over a period time.

Window functions aren't supported for MySQL versions earlier than 8 and MariaDB versions earlier than 10.2.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```

windowMin
(
    measure
    , [sort_order_field ASC/DESC, ...]
    , start_index
    , end_index
    , [ partition_field, ... ]

```

```
)
```

Arguments

measure

The aggregated metric that you want to get the average for, for example `sum( {Revenue} )`.

sort attribute

One or more aggregated fields, either measures or dimensions or both, that you want to sort the data by, separated by commas. You can either specify ascending (**ASC**) or descending (**DESC**) sort order.

Each field in the list is enclosed in `{ }` (curly braces), if it's more than one word. The entire list is enclosed in `[ ]` (square brackets).

start index

The start index is a positive integer, indicating n rows above the current row. The start index counts the available data points above the current row, rather than counting actual time periods. If your data is sparse (missing months or years, for example), adjust the indexes accordingly.

end index

The end index is a positive integer, indicating n rows below the current row. The end index counts the available data points below the current row, rather than counting actual time periods. If your data is sparse (missing months or years, for example), adjust the indexes accordingly.

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in `{ }` (curly braces), if it's more than one word. The entire list is enclosed in `[ ]` (square brackets).

Example


The following example calculates the trailing 12-month minimum of `sum(Revenue)`, partitioned by `SaleDate`. The calculation includes 12 rows above and 0 row below of the current row.

```

windowMin
(
    sum(Revenue),
    [SaleDate ASC],
    12,
    0
)

```

The following screenshot shows the results of this trailing 12-month example. The `sum(Revenue)` field is added to the chart to show the difference between the revenue and the trailing 12-month minimum revenue.



windowSum

The `windowSum` function calculates the sum of the aggregated measure in a custom window that is partitioned and sorted by specified attributes. Usually, you use custom window functions on a time series, where your visual shows a metric and a date field.

Window functions aren't supported for MySQL versions earlier than 8 and MariaDB versions earlier than 10.2.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
windowSum
(
    measure
    , [ sort_order_field ASC/DESC, ... ]
    , start_index
    , end_index
    , [ partition_field, ... ]
)
```

Arguments

measure

The aggregated metric that you want to get the sum for, for example `sum( {Revenue} )`.

For the engines MySQL, MariaDB, and Amazon Aurora with MySQL compatibility, the lookup index is limited to just 1. Window functions aren't supported for MySQL versions below 8 and MariaDB versions earlier than 10.2.

sort attribute

One or more aggregated fields, either measures or dimensions or both, that you want to sort the data by, separated by commas. You can either specify ascending (**ASC**) or descending (**DESC**) sort order.

Each field in the list is enclosed in {} (curly braces), if it's more than one word. The entire list is enclosed in [] (square brackets).

start index

The start index is a positive integer, indicating n rows above the current row. The start index counts the available data points above the current row, rather than counting actual time periods. If your data is sparse (missing months or years, for example), adjust the indexes accordingly.

end index

The end index is a positive integer, indicating n rows below the current row. The end index counts the available data points below the current row, rather than counting actual time

periods. If your data is sparse (missing months or years, for example), adjust the indexes accordingly.

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in {} (curly braces), if it's more than one word. The entire list is enclosed in [] (square brackets).

Example


The following example calculates the moving sum of `sum(Revenue)`, sorted by `SaleDate`. The calculation includes two rows above and one row ahead of the current row.

```
windowSum
(
    sum(Revenue),
    [SaleDate ASC],
    2,
    1
)
```

The following example show a trailing 12-month sum.

```
windowSum(sum(Revenue), [SaleDate ASC], 12, 0)
```

The following screenshot shows the results of this trailing 12-month sum example. The `sum(Revenue)` field is added to the chart to show the difference between the revenue and the trailing 12-month sum of revenue.



Joining data

You can use the join interface in Amazon Quick Sight to join objects from one or more data sources. By using Amazon Quick Sight to join the data, you can merge disparate data without duplicating the data from different sources.

Types of joined datasets

A join is performed between two Quick Sight *logical tables*, where each logical table contains information about how to fetch data. When editing a dataset in Quick Sight, the join diagram at the top half of the page shows each logical table as a rectangular block.

There are two different types of joined datasets in Quick Sight: same-source and cross-source. A dataset is considered same-source when it doesn't have any joins, or when all of the following conditions are met:

- If any of the logical tables refer to a Quick Sight data source:
 - All of the logical tables in this dataset must refer to the same Quick Sight data source. This doesn't apply if two separate Quick Sight data sources refer to the same underlying database.

It must be the exact same Quick Sight data source. For more information about using a single data source, see [Creating a dataset using an existing data source](#).

- If any of the logical tables refer to a Quick Sight dataset that is a parent dataset:
 - The parent dataset must use direct query.
 - The parent dataset must refer to the same Quick Sight data source.

If the above conditions aren't met, the dataset is considered a cross-source join.

Facts about joining datasets

Both same-source and cross-source dataset joins have the following limitations.

What's the maximum number of tables a joined dataset can contain?

All joined datasets can contain up to 32 tables.

How large can joined data be?

The maximum allowed size of a join is determined by the query mode and query engine that is used. The list below provides information about the different size limits for the tables to be joined. The size limit applies to all secondary tables combined. There are no join size limits for the primary table.

- **Same-source tables** – When tables originate from a single query data source, Quick Sight imposes no restrictions on the join size. This does not override join size limitations that the source query engine may have in place.
- **Cross-source datasets** – This type of join contains tables from different data sources that aren't stored in SPICE. For these types of joins, Quick Sight automatically identifies the largest table in the dataset. The combined size of all other secondary tables must be less than 1 GB.
- **Datasets stored in SPICE** – This type of join contains tables that are all ingested into SPICE. The combined size of all secondary tables in this join cannot exceed 20 GB.

For more information about SPICE dataset size calculations, see [Estimating the size of SPICE datasets](#).

Can a joined dataset use direct query?

Same-source datasets support direct query, assuming there are no other restrictions on using direct query. For example, S3 data sources don't support direct query, so a same-source S3 dataset must still use SPICE.

Cross-source datasets must use SPICE.

Can calculated fields be used in a join?

All joined datasets can use calculated fields, but calculated fields can't be used in any on-clauses.

Can geographical data be used in a join?

Same-source datasets support geographical data types, but geographical fields can't be used in any on-clauses.

Cross-source datasets don't support geographical data in any form.

For some examples of joining tables across data sources, see the [Joining across data sources on Amazon Quick Sight](#) post on the AWS Big Data Blog.

Creating a join

Use the following procedure to join tables to use in a dataset. Before you begin, import or connect to your data. You can create a join between any of the data sources supported by Amazon Quick Sight, except Internet of Things (IoT) data. For example, you can add comma-separated value (.csv) files, tables, views, SQL queries, or JSON objects in an Amazon S3 bucket.

To add one or more joins

1. Open the dataset that you want to work with.
2. (Optional) Before you get started, decide if you want to disable the autogenerated preview based on of a sample of your data. To turn that off, choose **Auto-preview** at top right. It's turned on by default.
3. If you haven't already chosen a query mode, choose **Query mode**.

Choose **SPICE** to store your dataset in [SPICE](#), or choose **Direct query** to pull live data every time. If your dataset contains one ore more manually uploaded file, your dataset is automatically stored in SPICE.

If you choose **SPICE**, the data is ingested into Quick Sight. Visuals that use the dataset run queries in SPICE, instead of on the database.

If you choose **Direct query**, the data isn't ingested into SPICE. Visuals that use the dataset run queries on the database, instead of in SPICE.

If you choose **Query mode**, make sure to set unique keys in the join, if applicable, to improve performance when loading visuals.

4. On the data preparation page, choose **Add data**.
5. In the **Add data page** that opens, choose one of the following options and complete the steps following:
 - Add data from a dataset:
 1. Choose **Dataset**.
 2. Select a dataset from the list.
 3. Choose **Select**.
 - Add data from a data source:
 1. Choose **Data source**.
 2. Select a data source from the list.
 3. Choose **Select**.
 4. Select a table from the list.
 5. Choose **Select**.
 - Create self-joins by adding a table multiple times. A counter appears after the name. An example is **Product**, **Product (2)**, and **Product (3)**. Field names in the **Fields** or **Filters** sections include the same counter so you can know which instance of the table a field came from.
 - Add a new file by choosing **Upload a file**, and then choose the file that you want to join.
6. (Optional) Choose **Use custom SQL** to open the query editor and write a query for a SQL data source.
7. (Optional) After you add data, interact with each table by choosing its menu icon. Rearrange the tables by dragging and dropping them.

An icon with red dots appears to indicate that you need to configure this join. Two red dots appear for joins that aren't yet configured. To create joins, choose the first join configuration icon.

8. (Optional) To change an existing join, reopen **Join configuration** by choosing the join icon between two tables.

The **Join Configuration** pane opens. On the join interface, specify the join type and the fields to use to join the tables.

9. At the bottom of the screen, you can see options to set a field in one table equal to a field in another table.

- In the **Join clauses** section, choose the join column for each table.

(Optional) If the tables that you selected join on multiple columns, choose **Add a new join clause**. Doing this adds another row to the join clauses, so you can specify the next set of columns to join. Repeat this process until you have identified all of the join columns for the two data objects.

10. In the **Join configuration** pane, choose the kind of join to apply. If the join fields are a unique key for one or both tables, enable the unique key setting. Unique keys only apply to direct queries, not to SPICE data.

For more information about joins, see [Join types](#).

11. Choose **Apply** to confirm your choices.

To cancel without making changes, choose **Cancel**.

12. The join icon in the workspace changes to show the new relationship.

13. (Optional) In the **Fields** section, you can use each field's menu to do one or more of the following:

- **Add a hierarchy** to a geospatial field.
- **Include** or **Exclude** the field.
- **Edit name & description** of the field.
- **Change data type**.
- **Add a calculation** (a calculated field).
- **Restrict access to only me**, so only you can see it. This can be helpful when you are adding fields to a dataset that's already in use.

14. (Optional) In the **Filters** section, you can add or edit filters. For more information, see [Filtering data in Amazon Quick Sight](#).

Join types

Amazon Quick Sight supports the following join types:

- Inner joins
- Left and right outer joins
- Full outer joins

Let's take a closer look at what these join types do with your data. For our example data, we're using the following tables named `widget` and `safety_rating`.

```
SELECT * FROM safety-rating
```

```
rating_id safety_rating
```

```
1      A+
2      A
3      A-
4      B+
5      B
```

```
SELECT * FROM WIDGET
```

```
widget_id    widget safety_rating_id
```

```
1      WidgetA  3
2      WidgetB  1
3      WidgetC  1
4      WidgetD  2
5      WidgetE
6      WidgetF  5
7      WidgetG
```

Inner joins

Use an inner join when you want to see only the data where there is a match between two tables. For example, suppose that you perform an inner join on the **safety-rating** and **widget** tables.

In the following result set, widgets without safety ratings are removed, and safety ratings without associated widgets are removed. Only the rows that match perfectly are included.

```
SELECT * FROM safety-rating
INNER JOIN widget
ON safety_rating.rating_id = widget.safety_rating_id
```

rating_id	safety_rating	widget_id	widget	safety_rating_id
3	A-	1	WidgetA	3
1	A+	2	WidgetB	1
1	A+	3	WidgetC	1
2	A	4	WidgetD	2
5	B	6	WidgetF	5

Left and right outer joins

These are also known as left or right outer joins. Use a left or right outer join when you want to see all the data from one table, and only the matching rows from the other table.

In a graphical interface, you can see which table is on the right or the left. In a SQL statement, the first table is considered to be on the left. Therefore, choosing a left outer join as opposed to a right outer join depends only on how the tables are laid out in your query tool.

For example, suppose that you perform a left outer join on `safety-rating` (the left table) and `widgets` (the right table). In this case, all `safety-rating` rows are returned, and only matching widget rows are returned. You can see blanks in the result set where there is no matching data.

```
SELECT * FROM safety-rating
LEFT OUTER JOIN widget
ON safety_rating.rating_id = widget.safety_rating_id
```

rating_id	safety_rating	widget_id	widget	safety_rating_id
1	A+	2	WidgetB	1
1	A+	3	WidgetC	1
2	A	4	WidgetD	2
3	A-	1	WidgetA	3
4	B+			
5	B	6	WidgetF	5

If you instead use a right outer join, call the tables in the same order so `safety-rating` is on the left and `widgets` is on the right. In this case, only matching `safety-rating` rows are returned,

and all widget rows are returned. You can see blanks in the result set where there is no matching data.

```
SELECT * FROM safety-rating
RIGHT OUTER JOIN widget
ON safety_rating.rating_id = widget.safety_rating_id
```

rating_id	safety_rating	widget_id	widget	safety_rating_id
3	A-	1	WidgetA	3
1	A+	2	WidgetB	1
1	A+	3	WidgetC	1
2	A	4	WidgetD	2
		5	WidgetE	
5	B	6	WidgetF	5
		7	WidgetG	

Full outer joins

These are sometimes called just outer joins, but this term can refer to either a left outer, right outer, or full outer join. To define the meaning, we use the complete name: full outer join.

Use a full outer join to see data that matches, plus data from both tables that doesn't match. This type of join includes all rows from both tables. For example, if you perform a full outer join on the safety-rating and widget tables, all rows are returned. The rows are aligned where they matched, and all extra data is included on separate rows. You can see blanks in the result set where there is no matching data.

```
SELECT * FROM safety-rating
FULL OUTER JOIN widget
ON safety_rating.rating_id = widget.safety_rating_id
```

rating_id	safety_rating	widget_id	widget	safety_rating_id
1	A+	2	WidgetB	1
1	A+	3	WidgetC	1
2	A	4	WidgetD	2
3	A-	1	WidgetA	3
4	B+			
5	B	6	WidgetF	5
		5	WidgetE	
		7	WidgetG	

Preparing data fields for analysis in Amazon Quick Sight

Before you start analyzing and visualizing your data, you can prepare the fields (columns) in your dataset for analysis. You can edit field names and descriptions, change the data type for fields, set up drill-down hierarchies for fields, and more.

Use the following topics to prepare fields in your dataset.

Topics

- [Editing field names and descriptions](#)
- [Setting fields as a dimensions or measures](#)
- [Changing a field data type](#)
- [Adding drill-downs to visual data in Quick Sight](#)
- [Selecting fields](#)
- [Organizing fields into folders in Amazon QuickSight](#)
- [Mapping and joining fields](#)

Editing field names and descriptions

You can change any field name and description from what is provided by the data source. If you change the name of a field used in a calculated field, make sure also to change it in the calculated field function. Otherwise, the function fails.

To change a field name or description

1. In the **Fields** pane of the data prep page, choose the three-dot icon on the field that you want to change. Then choose **Edit name & description**.
2. Enter the new name or description that you want to change, and choose **Apply**.

You can also change the name and description of a field on the data prep page. To do this, select the column header of the field that you want to change in the **Dataset** table in that page's lower half. Then make any changes there.

Setting fields as a dimensions or measures

In the **Field list** pane, dimension fields have blue icons and measure fields have green icons. *Dimensions* are text or date fields that can be items, like products, or attributes that are related to

measures. You can use dimensions to partition these items or attributes, like sales date for sales figures. *Measures* are numeric values that you use for measurement, comparison, and aggregation.

In some cases, Quick Sight interprets a field as a measure that you want to use it as a dimension (or the other way around). If so, you can change the setting for that field.

Changing a field's measure or dimension setting changes it for all visuals in the analysis that use that dataset. However, it doesn't change it in the dataset itself.

Changing a field's dimension or measure setting

Use the following procedure to change a field's dimension or measure setting

To change a field's dimension or measure setting

1. In the **Field list** pane, hover over the field that you want to change.
2. Choose the selector icon to the right of the field name, and then choose **Convert to dimension** or **Convert to measure** as appropriate.

Changing a field data type

When Quick Sight retrieves data, it assigns each field a data type based on the data in the field. The possible data types are as follows:

- **Date** – The date data type is used for date data in a supported format. For information about the date formats Quick Sight supports, see [Data source quotas](#).
- **Decimal** – The decimal data type is used for numeric data that requires one or more decimal places of precision, for example 18.23. The decimal data type supports values with up to four decimal places to the right of the decimal point. Values that have a higher scale than this are truncated to the fourth decimal place in two cases. One is when these values are displayed in data preparation or analyses, and one is when these values are imported into Quick Sight. For example, 13.00049 is truncated to 13.0004.
- **Geospatial** – The geospatial data type is used for geospatial data, for example longitude and latitude, or cities and countries.
- **Integer** – The int data type is used for numeric data that only contains integers, for example 39.
- **String** – The string data type is used for nondatetime alphanumeric data.

Quick Sight reads a small sample of rows in the column to determine the data type. The data type that occurs most in the small sample size is the suggested type. In some cases, there might be blank values (treated as strings by Quick Sight) in a column that contains mostly numbers. In these cases, it might be that the String data type is the most frequent type in the sample set of rows. You can manually modify the data type of the column to make it integer. Use the following procedures to learn how.

Changing a field data type during data prep

During data preparation, you can change the data type of any field from the data source. On the **Change data type** menu, you can change calculated fields that don't include aggregations to geospatial types. You can make other changes to the data type of a calculated field by modifying its expression directly. Quick Sight converts the field data according to the data type that you choose. Rows that contain data that is incompatible with that data type are skipped. For example, suppose that you convert the following field from String to Integer.

```
10020
36803
14267a
98457
78216b
```

All records containing alphabetic characters in that field are skipped, as shown following.

```
10020
36803
98457
```

If you have a database dataset with fields whose data types aren't supported by Quick Sight, use a SQL query during data preparation. Then use CAST or CONVERT commands (depending on what is supported by the source database) to change the field data types. For more information about adding a SQL query during data preparation, see [Using SQL to customize data](#). For more information about how different source data types are interpreted by Quick Sight, see [Supported data types from external data sources](#).

You might have numeric fields that act as dimensions rather than metrics, for example ZIP codes and most ID numbers. In these cases, it's helpful to give them a string data type during data preparation. Doing this lets Quick Sight understand that they are not useful for performing mathematical calculations and can only be aggregated with the Count function. For more

information about how Quick Sight uses dimensions and measures, see [Setting fields as a dimensions or measures](#).

In [SPICE](#), numbers converted from numeric into an integer are truncated by default. If you want to round your numbers instead, you can create a calculated field using the [round](#) function. To see whether numbers are rounded or truncated before they are ingested into SPICE, check your database engine.

To change a field data type during data prep

1. From the Quick Sight homepage, choose **Data** at left. In the **Data** tab, choose the dataset that you want, and then choose **Edit dataset**.
2. In the data preview pane, choose the data type icon under the field you want to change.
3. Choose the target data type. Only data types other than the one currently in use are listed.

Changing a field data type in an analysis

You can use the **Field list** pane, visual field wells, or on-visual editors to change numeric field data types within the context of an analysis. Numeric fields default to displaying as numbers, but you can choose to have them display as currency or as a percentage instead. You can't change the data types for string or date fields.

Changing a field's data type in an analysis changes it for all visuals in the analysis that use that dataset. However, it doesn't change it in the dataset itself.

Note

If you are working in a pivot table visual, applying a table calculation changes the data type of the cell values in some cases. This type of change occurs if the data type doesn't make sense with the applied calculation.

For example, suppose that you apply the Rank function to a numeric field that you modified to use a currency data type. In this case, the cell values display as numbers rather than currency. Similarly, if you apply the Percent difference function instead, the cell values display as percentages rather than currency.

To change a field's data type

1. Choose one of the following options:

- In the **Field list** pane, hover over the numeric field that you want to change. Then choose the selector icon to the right of the field name.
 - On any visual that contains an on-visual editor associated with the numeric field that you want to change, choose that on-visual editor.
 - Expand the **Field wells** pane, and then choose the field well associated with the numeric field that you want to change.
2. Choose **Show as**, and then choose **Number**, **Currency**, or **Percent**.

Adding drill-downs to visual data in Quick Sight

All visual types except pivot tables offer the ability to create a hierarchy of fields for a visual element. The hierarchy lets you drill down to see data at different levels of the hierarchy. For example, you can associate the country, state, and city fields with the x-axis on a bar chart. Then, you can drill down or up to see data at each of those levels. As you drill down each level, the data displayed is refined by the value in the field you drill down on. For example, if you drill down on the state of California, you see data on all of the cities in California.

The field wells you can use to create drill-downs varies by visual type. Refer to the topic on each visual type to learn more about its drill-down support.

Drill-down functionality is added automatically for dates when you associate a date field with the drill-down field well of a visual. In this case, you can always drill up and down through the levels of date granularity. Drill-down functionality is also added automatically for geospatial groupings, after you define these in the dataset.

Use the following table to identify the field wells/on-visual editors that support drill-down for each visual type.

Visual type	Field well or on-visual editor
Bar charts (all horizontal)	Y axis and Group/Color
Bar charts (all vertical)	X axis and Group/Color
Combo charts (all)	X axis and Group/Color
Geospatial charts	Geospatial and Color

Visual type	Field well or on-visual editor
Heat map	Rows and Columns
KPIs	Trend Group
Line charts (all)	X axis and Color
Pie chart	Group/Color
Scatter plot	Group/Color
Tree map	Group by

Important

Drill-downs are not supported for tables or pivot tables.

Adding a drill-down

Use the following procedure to add drill-down levels to a visual.

To add drill-down levels to a visual

1. On the analysis page, choose the visual that you want to add drill-downs to.
2. Drag a field item into a **Field well**.
3. If your dataset has a defined hierarchy, you can drag the entire hierarchy into the field well as one. An example is geospatial or coordinate data. In this case, you don't need to follow the remaining steps.

If you don't have a predefined hierarchy, you can create one in your analysis, as described in the remaining steps.

4. Drag a field that you want to use in the drill-down hierarchy to an appropriate field well, depending on the visual type. Make sure that the label for the dragged field says **Add drill-down layer**. Position the dragged field above or below the existing field based on where you want it to be in the hierarchy you're creating.

5. Continue until you have added all of the levels of hierarchy that you want. To remove a field from the hierarchy, choose the field, and then choose **Remove**.
6. To drill down or up to see data at a different level of the hierarchy, choose an element on the visual (like a line or bar), and then choose **Drill down to <lower level>** or **Drill up to <higher level>**. In this example, from the `car-make` level you can drill down to `car-model` to see data at that level. If you drill down to `car-model` from the **Ford** `car-make`, you see only `car-models` in that `car-make`.

After you drill down to the `car-model` level, you can then drill down further to see `make-year` data, or go back up to `car-make`. If you drill down to `make-year` from the bar representing **Ranger**, you see only years for that model of car.

Selecting fields

When you prepare data, you can select one or more fields to perform an action on them, such as excluding them or adding them to a folder.

To select one or more fields in the data preparation pane, click or tap the field or fields in the **Fields** pane at left. You can then choose the field menu (the three dots) to the right of the field name and choose an action to take. The action is performed on all selected fields.

You can select or deselect all fields at once by choosing either **All** or **None** at the top of the **Fields** pane.

If you edit a dataset and exclude a field that is used in a visual, that visual breaks. You can fix it the next time you open that analysis.

Searching for fields

If you have a long field list in the **Fields** pane, you can search to locate a specific field by entering a search term for **Search fields**. Any field whose name contains the search term is shown.

Search is case-insensitive and wildcards are not supported. Choose the cancel icon (X) to the right of the search box to return to viewing all fields.

Organizing fields into folders in Amazon QuickSight

When prepping your data in Quick Sight, you can use folders to organize your fields for multiple authors across your enterprise. Arranging fields into folders and subfolders can make it easier for authors to find and understand fields in your dataset.

You can create folders while preparing your dataset, or when editing a dataset. For more information about creating a new dataset and preparing it, see [Creating datasets](#). For more information about opening an existing dataset for data preparation, see [Editing datasets](#).

While performing an analysis, authors can expand and collapse folders, search for specific fields within folders, and see your descriptions of folders on the folder menu. Folders appear at the top of the **Fields** pane in alphabetical order.

Creating a folder

Use the following procedure to create a new folder in the **Fields** pane.

To create a new folder

1. On the data preparation page, in the **Fields** pane, select the three-dot icon, and choose **Add to folder**.

To select more than one field at a time, press the Ctrl key while you select (Command key on Mac).

2. On the **Add to folder** page that appears, choose **Create a new folder** and enter a name for the new folder.
3. Choose **Apply**.

The folder appears at the top of the **Fields** pane with the fields that you chose inside it. Fields inside folders are arranged in alphabetical order.

Creating a subfolder

To further organize your data fields in the **Fields** pane, you can create subfolders within parent folders.

To create a subfolder

1. On the data preparation page, in the **Fields** pane, select the field menu for a field already in a folder and choose **Move to folder**.
2. On the **Move to folder** page that appears, choose **Create a new folder** and enter a name for the new folder.
3. Choose **Apply**.

The subfolder appears within the parent folder at the top of the list of fields. Subfolders are arranged in alphabetical order.

Adding fields to an existing Folder

Use the following procedure to add fields to an existing folder in the **Fields** pane.

To add one or more fields to a folder

1. On the data preparation page, in the **Fields** pane, select the fields that you want to add to a folder.

To select more than one field at a time, press the Ctrl key while you select (Command key on Mac).

2. On the field menu, choose **Add to folder**.
3. On the **Add to folder** page that appears, choose a folder for **Existing folder**.
4. Choose **Apply**.

The field or fields are added to the folder.

Moving fields between folders

Use the following procedure to move fields between folders in the **Fields** pane.

To move fields between folders

1. On the data preparation page, in the **Fields** pane, select the fields that you want to move to another folder.

To select more than one field at a time, press the Ctrl key while you select (Command key on Mac).

2. On the field menu, choose **Move to folder**.
3. On the **Move to folder** page that appears, choose a folder for **Existing folder**.
4. Choose **Apply**.

Removing fields from a folder

Use the following procedure to remove fields from a folder in the **Fields** pane. Removing a field from a folder doesn't delete the field.

To remove fields from a folder

1. On the data preparation page, in the **Fields** pane, select the fields that you want to remove.
2. On the field menu, choose **Remove from folder**.

The fields that you selected are removed from the folder and placed back in the list of fields in alphabetical order.

Editing a folder name and adding a folder description

You can edit the name or add a description of a folder to provide context about the data fields inside it. The folder name appears in the **Fields** pane. While performing an analysis, authors can read your folder's description when they select the folder menu in the **Fields** pane.

To edit a folder name or edit or add a description for a folder

1. On the data preparation page, in the **Fields** pane, select the folder menu for the folder that you want to edit and choose **Edit name & description**.
2. On the **Edit folder** page that appears, do the following:
 - For **Name**, enter a name for the folder.
 - For **Description**, enter a description of the folder.
3. Choose **Apply**.

Moving folders

You can move folders and subfolders to new or existing folders in the **Fields** pane.

To move a folder

1. On the data preparation page, in the **Fields** pane, choose **Move folder** on the folder menu.
2. On the **Move folder** page that appears, do one of the following:
 - Choose **Create a new folder** and enter a name for the folder.
 - For **Existing folder**, choose a folder.
3. Choose **Apply**.

The folder appears within the folder that you chose in the **Fields** pane.

Removing folders from the fields pane

Use the following procedure to remove a folder from the **Fields** pane.

To remove a folder

1. On the data preparation page, in the **Fields** pane, choose **Remove folder** on the folder menu.
2. On the **Remove folder?** page that appears, choose **Remove**.

The folder is removed from the **Fields** pane. Any fields that were in the folder are placed back in the list of fields in alphabetical order. Removing folders doesn't exclude fields from view or delete fields from the dataset.

Mapping and joining fields

When you are using different datasets together in Quick Sight, you can simplify the process of mapping fields or joining tables during the data preparation stage. You should already be verifying that your fields have the correct data type and an appropriate field name. However, if you already know which datasets are going to be used together, you can take a couple of extra steps to make your work easier later on.

Mapping fields

Quick Sight can automatically map fields between datasets in the same analysis. The following tips can help make it easier for Quick Sight to automatically map fields between datasets, for example if you are creating a filter action across datasets:

- Matching field names – Field names must match exactly, with no differences in case, spacing, or punctuation. You can rename fields that describe the same data, so an automatic mapping is accurate.
- Matching data types – Fields must have the same data type for automatic mapping. You can change the data types while you are preparing the data. This step also gives you the opportunity to discover whether you need to filter out any data that isn't the correct data type.
- Using calculated fields – You can use calculated fields to create a matching field, and give it the correct name and data type for automatic mapping.

Note

After an automatic mapping exists, you can rename a field without breaking the field mapping. However, if you change the data type, the mapping is broken.

For more information on field mapping for filter actions across datasets, see [Creating and editing custom actions in Amazon Quick Sight](#).

Joining fields

You can create joins between data from different data sources, including files or databases. The following tips can help make it easier for you to join data from different files or data sources:

- **Similar field names** – It is simpler to join fields when you can see what should match; for example, **Order ID** and **order-id** seem as if they should be the same. But if one is a work order, and the other is a purchase order, then the fields are probably different data. If possible, make sure that the files and tables that you want to join have field names making it clear what data they contain.
- **Matching data types** – Fields must have the same data type before you can join on them. Make sure that the files and tables that you want to join having matching data types in join fields. You can't use a calculated field for a join. Also, you can't join two existing datasets. You create the joined dataset by directly accessing the source data.

For more information on joining data across data sources, see [Joining data](#).

Filtering data in Amazon Quick Sight

You can use filters to refine the data in a dataset or an analysis. For example, you can create a filter on a region field that excludes data from a particular region in a dataset. You can also add a filter to an analysis, such as a filter on the range of dates that you want to include in any visuals in your analysis.

When you create a filter in a dataset, that filter applies to the entire dataset. Any analyses and subsequent dashboards created from that dataset contains the filter. If someone creates a dataset from your dataset, the filter also is in the new dataset.

When you create a filter in an analysis, that filter only applies to that analysis and any dashboards you publish from it. If someone duplicates your analysis, the filter persists in the new analysis. In

analyses, you can scope filters to a single visual, some visuals, all visuals that use this dataset, or all applicable visuals.

Also, when you create filters in an analysis, you can add a filter control to your dashboard. For more information about filter controls, see [Adding filter controls to analysis sheets](#).

Each filter you create applies only to a single field. You can apply filters to both regular and calculated fields.

There are several types of filters you can add to datasets and analyses. For more information about the types of filters you can add, and some of their options, see [Filter types in Amazon Quick](#).

If you create multiple filters, all top-level filters apply together using AND. If you group filters by adding them inside a top-level filter, the filters in the group apply using OR.

Amazon Quick Sight applies all of the enabled filters to the field. For example, suppose that there is one filter of `state = WA` and another filter of `sales >= 500`. Then the dataset or analysis only contains records that meet both of those criteria. If you disable one of these, only one filter applies.

Take care that multiple filters applied to the same field aren't mutually exclusive.

Use the following sections to learn how to view, add, edit, and delete filters.

Topics

- [Viewing existing filters](#)
- [Adding filters](#)
- [Cross-sheet filters and controls](#)
- [Filter types in Amazon Quick](#)
- [Adding filter controls to analysis sheets](#)
- [Editing filters](#)
- [Enabling or disabling filters](#)
- [Deleting filters](#)

Viewing existing filters

When you edit a dataset or open an analysis, you can view any existing filters that were created. Use the following procedures to learn how.

Viewing filters in datasets

1. Open the [Quick console](#).
2. From the Quick homepage, choose **Data** at left.
3. In the **Datasets** tab, choose the dataset that you want, and then choose **Edit dataset**.
4. On the data preparation page that opens, choose **Filters** at lower left to expand the **Filters** section.

Any filters that are applied to the dataset appear here. If a single field has multiple filters, they are grouped together. They display in order of create date, with the oldest filter on top.

Viewing filters in analyses

Use the following procedure to view filters in analyses.

To view a filter in an analysis

1. From the Quick homepage, choose **Analyses**.
2. On the **Analyses** page, choose the analysis that you want to work with.
3. In the analysis, choose the **Filter** icon to open the **Filters** pane.

Any filters applied to the analysis appear here.

The way that a filter is scoped is listed at the bottom of each filter. For more information about scoping filters, see [Adding filters](#).

Adding filters

You can add filters to a dataset or an analysis. Use the following procedures to learn how.

Adding filters to datasets

Use the following procedure to add filters to datasets.

To add a filter to a dataset

1. Open the [Quick console](#).
2. From the Quick homepage, choose **Data** at left.

3. In the **Datasets** tab, choose the dataset that you want, and then choose **Edit dataset**.
4. On the data preparation page that opens, choose **Add filter** at lower left, and then choose a field that you want to filter.

The filter is added to the **Filters** pane.

5. Choose the new filter in the pane to configure the filter. Or you can choose the three dots to the right of the new filter and choose **Edit**.

Depending on the data type of the field, your options for configuring the filter vary. For more information about the types of filters that you can create and their configurations, see [Filter types in Amazon Quick](#).

6. When finished, choose **Apply**.

Note

The data preview shows you the results of your combined filters only as they apply to the first 1,000 rows. If all of the first 1,000 rows are filtered out, then no rows show in the preview. This effect occurs even when rows after the first 1,000 aren't filtered out.

Adding filters in analyses

Use the following procedure to add filters to analyses.

To add a filter to an analysis

1. Open the [Quick console](#).
2. From the Quick homepage, choose **Analyses**.
3. On the **Analyses** page, choose the analysis that you want to work with.
4. In the analysis, choose the **Filter** icon to open the **Filters** pane, and then choose **ADD**.
5. Choose the new filter in the pane to configure it. Or you can choose the three dots to the right of the new filter and choose **Edit**.
6. In the **Edit filter** pane that opens, for **Applied to**, choose one of the following options.
 - **Single visual** – The filter applies to the selected item only.
 - **Single sheet** – The filter applies to a single sheet.
 - **Cross sheet** – The filter applies to multiple sheets in the dataset.

Depending on the data type of the field, your remaining options for configuring the filter vary. For more information about the types of filters you can create and their configurations, see [Filter types in Amazon Quick](#).

Cross-sheet filters and controls

Cross-sheet filters and controls are filters that are scoped to either your entire analysis or dashboard or multiple sheets within your analysis and dashboard.

Filters

Creating a Cross-Sheet Filter

1. Once you have [Added a filter](#), you update the scope of the filter to cross-sheet. By default, this applies to all of the the sheets in your analysis.
2. If the **Apply cross-datasets** box is checked, then the filter will be applied to all visuals from up to 100 different datasets that are applicable to all sheets in the filter scope.
3. If you want to customize the sheets that it is applied to, then choose the Cross-sheet icon. You can then view the sheets the filter is currently applied to or toggle on the custom select sheets.
4. When you enable **Custom select sheets**, you can select which sheets to apply the filter to.
5. Follow the steps at [Editing filters in analyses](#). Your changes will be applied to all of the filters for all of the sheets you have selected. This includes newly added sheets if the filter is scoped to your entire analysis.

Removing a Cross-Sheet Filter

Deleting

If you have no controls created from these filters, see [Deleting filters in analyses](#).

If you have controls created then:


1. Follow the instructions at [Deleting filters in analyses](#).
2. If you choose **Delete Filter and Controls**, the controls will be deleted from all pages. This may impact the layout of your analysis. Alternatively, you can remove these controls individually.

Downscoping

If you want to remove a cross-sheet filter, you can also do this by changing the filter scope:

1. Follow the instructions at [Editing filters in analyses](#) to get to the filter.
2. One of the edits you can make is changing the scope. You can switch to **Single sheet** or **Single visual**. You can also remove a sheet from the Cross-sheet selection.

Or the custom sheet selection:



3. If there are controls, you will see a modal to warn you that you will be bulk-removing controls from any of the sheets where the filter no longer applies and this can impact your layout. You can also remove the controls individually. For more information, see [Removing a Cross-Sheet Control](#).

4. If you add controls to the **Top of all sheets in filter scope** then new sheets will by default be added with this new control if the filter is scoped to your entire analysis.

Controls

Creating a Cross-Sheet Control

New filter control

1. Create a cross-sheet filter. For more information, see [Filters](#).
2. From the three-dot menu, you can see an option that says **Add control**. Hovering over this, you will see three options:
 - **Top of all sheets in filter scope**
 - **Top of this sheet**
 - **Inside this sheet**

If you want to add to multiple-sheets within the sheets themselves, you can do that sheet-by-sheet. Or you can add to the top and then use the option on each control to **Move to sheet**. For more information, see [Editing a Cross-Sheet Control](#).

Increasing Scope of Existing Control

1. Navigate to the existing filter in the analysis
2. Change the scope of what sheets this filter is **Applied to to Cross-sheet**.
3. If there is already a control created from the filter, you will see a modal, which if you check the box will bulk-add controls to the top of all the sheets in the filter scope. This will not impact the position of the already created control if it is on the sheet.

Editing a Cross-Sheet Control

1. Go to the cross-sheet control and select the three-dot menu if the control is pinned to the top or the edit pencil icon if the control is on the sheet. You will be presented with the following options:
 - **Go to filter** (which directs you to the cross-sheet filter for you to edit or review)
 - **Move to sheet** (which moves the control into the analysis pane)

- **Reset**
 - **Refresh**
 - **Edit**
 - **Remove**
2. Choose **Edit**. This brings up the **Format Control** pane on the right side of your analysis.
 3. You can then edit your control. The top section labeled **Cross-sheet settings** will apply to all controls, whereas any settings outside of this section are not applicable to all controls and only to the specific control you're editing. For instance, **Relevant value** is not a cross-sheet control setting.
 4. You can also see the sheets that this control is on as well as the location (Top or Sheet) that the control is on for each sheet. You can do this by choosing **Sheets(8)**.

Removing a Cross-Sheet Control

You can remove controls in two places. First, from the control:

1. Go to the cross-sheet control and select the three-dot menu if the control is pinned to the top or the edit pencil icon if the control is on the sheet. You will be presented with the following options:
 - **Go to filter** (which directs you to the cross-sheet filter for you to edit or review)
 - **Move to sheet** (which moves the control into the analysis pane)
 - **Reset**
 - **Refresh**
 - **Edit**
 - **Remove**
2. Choose **Remove**

Second, you can remove controls from the filter:

1. Choose the three-dot menu on the cross-sheet filter that the cross-sheet controls are created from. You will see that instead of an option to **Add control** there is now an option to **Manage control**.
2. Hover over **Manage control**. You will be presented with the following options:

- **Move inside this sheet**
- **Top of this sheet**

These options are applicable to just the control on the sheet, depending on where the current control is. If you don't have controls on all of the sheets within the filter scope, you will get the option to **Add to top of all sheets in filter scope**. This will not move sheet controls to the top of the sheet if you have already added them to the sheet in the analysis. You will also get the option to **Remove from this sheet** or **Remove from all sheets**.

Filter types in Amazon Quick

You can create several different types of filters in Quick. The type of filter you create mostly depends on the data type of the field that you want to filter.

In datasets, you can create the following types of filters:

- Text filters
- Numeric filters
- Date filters

In analyses, you can create the same types of filters as you can in datasets. You can also create:

- Group filters with and/or operators
- Cascading filters
- Nested filters

Use the following sections to learn more about each type of filter you can create and some of their options.

Topics

- [Adding text filters](#)
- [Adding nested filters](#)
- [Adding numeric filters](#)
- [Adding date filters](#)

- [Adding filter conditions \(group filters\) with AND and OR operators](#)
- [Creating cascading filters](#)

Adding text filters

When you add a filter using a text field, you can create the following types of text filters:

- **Filter list (Analyses only)** – This option creates a filter that you can use to select one or more field values to include or exclude from all the available values in the field. For more information about creating this type of text filter, see [Filtering text field values by a list \(analyses only\)](#).
- **Custom filter list** – With this option, you can enter one or more field values to filter on, and whether you want to include or exclude records that contain those values. The values that you enter must match the actual field values exactly for the filter to be applied to a given record. For more information about creating this type of text filter, see [Filtering text field values by a custom list](#).
- **Custom filter** – With this option, you enter a single value that the field value must match in some way. You can specify that the field value must equal, not equal, starts with, ends with, contains, or does not contain the value you specify. If you choose an equal comparison, the specified value and actual field value must match exactly in order for the filter to be applied to a given record. For more information about creating this type of text filter, see [Filtering a single text field value](#).
- **Top and bottom filter (Analyses only)** – You can use this option to show the top or bottom *n* value of one field ranked by the values in another field. For example, you might show the top five salespeople based on revenue. You can also use a parameter to allow dashboard users to dynamically choose how many top or bottom ranking values to show. For more information about creating top and bottom filters, see [Filtering a text field by a top or bottom value \(analyses only\)](#).

Filtering text field values by a list (analyses only)

In analyses, you can filter a text field by selecting values to include or exclude from a list of all value in the field.

To filter a text field by including and excluding values

1. Create a new filter using a text field. For more information about creating filters, see [Adding filters](#).

2. In the **Filters** pane, choose the new filter to expand it.
3. For **Filter type**, choose **Filter list**.
4. For **Filter condition**, choose **Include** or **Exclude**.
5. Choose the field values that you want to filter on. To do this, select the check box in front of each value.

If there are too many values to choose from, enter a search term into the box above the checklist and choose **Search**. Search terms are case-insensitive and wildcards aren't supported. Any field value that contains the search term is returned. For example, searching on L returns al, AL, la, and LA.

The values display alphabetically in the control, unless there are more than 1,000 distinct values. Then the control displays a search box instead. Each time that you search for the value that you want to use, it starts a new query. If the results contain more than 1,000 values, you can scroll through the values with pagination.

6. When finished, choose **Apply**.

Filtering text field values by a custom list

You can specify one or more field values to filter on, and whether you want to include or exclude records that contain those values. The specified value and actual field value must match exactly for the filter to be applied to a given record.

To filter text field values by a custom list

1. Create a new filter using a text field. For more information about creating filters, see [Adding filters](#).
2. In the **Filters** pane, choose the new filter to expand it.
3. For **Filter type**, choose **Custom filter list**.
4. For **Filter condition**, choose **Include** or **Exclude**.
5. For **List**, enter a value in the text box. The value must match an existing field value exactly.
6. (Optional) To add additional values, enter them in the text box, one per line.
7. For **Null options** choose **Exclude nulls**, **Include nulls**, or **Nulls only**.
8. When finished, choose **Apply**.

Filtering a single text field value

With the **Custom filter** filter type, you specify a single value that the field value must equal or not equal, or must match partially. If you choose an equal comparison, the specified value and actual field value must match exactly for the filter to be applied to a given record.

To filter a text field by a single value

1. Create a new filter using a text field. For more information about creating filters, see [Adding filters](#).
2. In the **Filters** pane, choose the new filter to expand it.
3. For **Filter type**, choose **Custom filter**.
4. For **Filter condition**, choose one of the following:
 - **Equals** – When you choose this option, the values included or excluded in the field must match the value that you enter exactly.
 - **Does not equal** – When you choose this option, the values included or excluded in the field must match the value that you enter exactly.
 - **Starts with** – When you choose this option, the values included or excluded in the field must start with the value that you enter.
 - **Ends with** – When you choose this option, the values included or excluded in the field must start with the value that you enter.
 - **Contains** – When you choose this option, the values included or excluded in the field must contain the whole value that you enter.
 - **Does not contain** – When you choose this option, the values included or excluded in the field must not contain any part of the value that you enter.

Note

Comparison types are case-sensitive.

5. Do one of the following:
 - For **Value**, enter a literal value.
 - Select **Use parameters** to use an existing parameter, and then choose a parameter from the list.

For parameters to appear in this list, create your parameters first. Usually, you create a parameter, add a control for it, and then add a filter for it. For more information, see [Parameters in Amazon Quick](#).

The values display alphabetically in the control, unless there are more than 1,000 distinct values. Then the control displays a search box instead. Each time that you search for the value that you want to use, it starts a new query. If the results contain more than 1,000 values, you can scroll through the values with pagination.

6. For **Null options** choose **Exclude nulls**, **Include nulls**, or **Nulls only**.
7. When finished, choose **Apply**.

Filtering a text field by a top or bottom value (analyses only)

You can use a **Top and bottom filter** to show the top or bottom n value of one field ranked by the values in another field. For example, you might show the top five salespeople based on revenue. You can also use a parameter to allow dashboard users to dynamically choose how many top or bottom ranking values to show.

To create a top and bottom text filter

1. Create a new filter using a text field. For more information about creating filters, see [Adding filters](#).
2. In the **Filters** pane, choose the new filter to expand it.
3. For **Filter type**, choose **Top and bottom filter**.
4. Choose **Top** or **Bottom**.
5. For **Show top** integer (or **Show bottom** integer), do one of the following:
 - Enter the number of top or bottom items to show.
 - To use a parameter for the number of top or bottom items to show, select **Use parameters**. Then choose an existing integer parameter.

For example, let's say that you want to show the top three salespersons by default. However, you want the dashboard viewer to be able to choose whether to show 1–10 top salespersons. In this case, take the following actions:

- Create an integer parameter with a default value.

- To link the number of displayed items to a parameter control, create a control for the integer parameter. Then you make the control a slider with a step size of 1, a minimum value of 1, and a maximum value of 10.
 - To make the control work, link it to a filter by creating a top and bottom filter on Salesperson by Weighted Revenue, enable **Use parameters**, and choose your integer parameter.
6. For **By**, choose a field to base the ranking on. If you want to show the top five salespeople per revenue, choose the revenue field. You can also set the aggregate that you want to perform on the field.
 7. (Optional) Choose **Tie breaker** and then choose another field to add one or more aggregations as tie breakers. This is useful, in the case of this example, when there are more than five results returned for the top five salespeople per revenue. This situation can happen if multiple salespeople have the same revenue amount.

To remove a tie breaker, use the delete icon.

8. When finished, choose **Apply**.

Adding nested filters

Nested filters are advanced filters that can be added to a Quick analysis. A Nested filter filters a field using a subset of data defined by another field in that same dataset. This allows authors to show additional contextual data without the need to filter data out if the data point doesn't meet an initial condition.

Nested filters function similarly to a correlated subquery in SQL or a market basket analysis. For example, say you want to perform a market basket analysis on your sales data. You can use nested filters to find the sales quantity by product for customers who have or have not purchased a specific product. You can also use nested filters to identify groups of customers that did not purchase a selected product or who only purchased a specific list of products.

Nested filters can only be added at the analysis level. You can't add a nested filter to a dataset.

Use the procedure below to add a nested filter to a Quick analysis.

1. Open the [Quick console](#).
2. Choose **Analyses**, and then choose the analysis that you want to add a nested filter to.

3. Create a new filter on the text field that you want to filter on. For more information about creating a filter, see [Adding filters in analyses](#).
4. After you create the new filter, locate the new filter in the **Filters** pane. Choose the ellipsis (three dots) next to the new filter, and then choose **Edit filter**. Alternatively, choose the filter entity in the **Filters** pane to open the **Edit filter** pane.
5. The **Edit filter** pane opens. Open the **Filter type** dropdown menu, navigate to the **Advanced filter** section, and then choose **Nested filter**.
6. For **Qualifying condition**, choose **Include** or **Exclude**. The *qualifying condition* allows you to run a not in the set query on the data in your analysis. In our sales example above, the qualifying condition determines if the filter returns a list of customers who did buy the specific product or a list of customers who did not buy the product.
7. For **Nested field**, choose the text field that you want to filter data with. The nested field cannot be the same as the primary field selected in step 3. Category fields are the only supported field type for the inner filter.
8. For **Nested filter type**, choose the filter type that you want. The filter type that you choose determines the final configuration steps for the nested filter. Available filter types and information about their configuration can be found in the list below.
 - [Filter list](#)
 - [Custom filter list](#)
 - [Custom filter](#)

Adding numeric filters

Fields with decimal or int data types are considered numeric fields. You create filters on numeric fields by specifying a comparison type, for example **Greater than** or **Between**, and a comparison value or values as appropriate to the comparison type. Comparison values must be positive integers and can't contain commas.

You can use the following comparison types in numeric filters:

- Equals
- Does not equal
- Greater than
- Greater than or equal to

- Less than
- Less than or equal to
- Between

Note

To use a top and bottom filter for numeric data (analyses only), first change the field from a measure to a dimension. Doing this converts the data to text. Then you can use a text filter. For more information, see [Adding text filters](#).

In analyses, for datasets based on database queries, you can also optionally apply an aggregate function to the comparison value or values, for example **Sum** or **Average**.

You can use the following aggregate functions in numeric filters:

- Average
- Count
- Count distinct
- Max
- Median
- Min
- Percentile
- Standard deviation
- Standard deviation - population
- Sum
- Variance
- Variance - population

Creating numeric filters

Use the following procedure to create a numeric field filter.

To create a numeric field filter

1. Create a new filter using a text field. For more information about creating filters, see [Adding filters](#).
2. In the **Filters** pane, choose the new filter to expand it.
3. (Optional) For **Aggregation**, choose an aggregation. No aggregation is applied by default. This option is available only when creating numeric filters in an analysis.
4. For **Filter condition**, choose a comparison type.
5. Do one of the following:

- If you chose a comparison type other than **Between**, enter a comparison value.

If you chose a comparison type of **Between**, enter the beginning of the value range in **Minimum value** and the end of the value range in **Maximum value**.

- (Analyses only) To use an existing parameter, enable **Use parameters**, then choose your parameter from the list.

For parameters to appear in this list, create your parameters first. Usually, you create a parameter, add a control for it, and then add a filter for it. For more information, see [Parameters in Amazon Quick](#). The values display alphabetically in the control, unless there are more than 1,000 distinct values. Then the control displays a search box instead. Each time you search for the value that you want to use, it initiates a new query. If the results contain more than 1,000 values, you can scroll through the values with pagination.

6. (Analyses only) For **Null options** choose **Exclude nulls**, **Include nulls**, or **Nulls only**.
7. When finished, choose **Apply**.

Adding date filters

You create filters on date fields by selecting the filter conditions and date values that you want to use. There are three filter types for dates:

- **Range** – A series of dates based on a time range and comparison type. You can filter records based on whether the date field value is before or after a specified date, or within a date range. You enter date values in the format MM/DD/YYYY. You can use the following comparison types:
 - **Between** – Between a start date and an end date
 - **After** – After a specified date

- **Before** – Before a specified date
- **Equals** – On a specified date

For each comparison type, you can alternatively choose a rolling date relative to a period or dataset value.

- **Relative** (analyses only) – A series of date and time elements based on the current date. You can filter records based on the current date and your selected unit of measure (UOM). Date filter units include years, quarters, months, weeks, days, hours, and minutes. You can exclude current period, add support for Next N filters similar to Last N with an added capability to allow for Anchor date. You can use the following comparison types:
 - **Previous** – The previous UOM—for example, the previous year.
 - **This** – This UOM, which includes all dates and times that fall within the select UOM, even if they occur in the future.
 - **To date or up to now** – UOM to date, or UOM up to now. The displayed phrase adapts to the UOM that you choose. However, in all cases this option filters out data that is not between the beginning of the current UOM and the current moment.
 - **Last n** – The last specified number of the given UOM, which includes all of this UOM and all of the last $n - 1$ UOM. For example, let's say today is May 10, 2017. You choose to use *years* as your UOM, and set Last n years to 3. The filtered data includes data for all of 2017, plus all of 2016, and all of 2015. If you have any data for the future dates of the current year (2017 in this example), these records are included in your dataset.
- **Top and bottom** (analyses only) – A number of date entries ranked by another field. You can show the top or bottom n for the type of date or time UOM you choose, based on values in another field. For example, you can choose to show the top 5 sales days based on revenue.

Comparisons are applied inclusive to the date specified. For example, if you apply the filter **Before 1/1/16**, the records returned include all rows with date values through 1/1/16 23:59:59. If you don't want to include the date specified, you can clear the option to **Include this date**. If you want to omit a time range, you can use the **Exclude the last N periods** option to specify the number and type of time periods (minutes, days, and so on) to filter out.

You can also choose to include or exclude nulls, or exclusively show rows that contain nulls in this field. If you pass in a null date parameter (one without a default value), it doesn't filter the data until you provide a value.

Note

If a column or attribute has no time zone information, then the client query engine sets the default interpretation of that date-time data. For example, suppose that a column contains a timestamp, rather than a timestamptz, and you are in a different time zone than the data's origin. In this case, the engine can render the timestamp differently than you expect. Amazon Quick and [SPICE](#) both use Universal Coordinated Time (UTC) times.

Use the following sections to learn how to create date filters in datasets and analyses.

Creating date filters in datasets

Use the following procedure to create a range filter for a date field in a dataset.

To create a range filter for a date field in a dataset

1. Create a new filter using a text field. For more information about creating filters, see [Adding filters](#).
2. In the **Filters** pane, choose the new filter to expand it.
3. For **Condition**, choose a comparison type: **Between**, **After**, or **Before**.

To use **Between** as a comparison, choose **Start date** and **End date** and choose dates from the date picker controls that appear.

You can choose if you want to include either or both the start and end dates in the range by selecting **Include start date** or **Include end date**.

To use **Before** or **After** comparisons, enter a date or choose the date field to bring up the date picker control and choose a date instead. You can include this date (the one you chose), to exclude the last N time periods, and specify how to handle nulls.

4. For **Time granularity**, choose **Day**, **Hour**, **Minute**, or **Second**.
5. When finished, choose **Apply**.

Creating date filters in analyses

You can create date filters in analyses as described following.

Creating range date filters in analyses

Use the following procedure to create a range filter for a date field in an analysis.

To create a range filter for a date field in an analysis

1. Create a new filter using a text field. For more information about creating filters, see [Adding filters](#).
2. In the **Filters** pane, choose the new filter to expand it.
3. For **Filter type**, choose **Date & time range**.
4. For **Condition**, choose a comparison type: **Between**, **After**, **Before**, or **Equals**.

To use **Between** as a comparison, choose **Start date** and **End date** and choose dates from the date picker controls that appear.

You can choose to include either or both the start and end dates in the range by selecting **Include start date** or **Include end date**.

To use a **Before**, **After**, or **Equals** comparison, enter a date or choose the date field to bring up the date picker control and choose a date instead. You can include this date (the one you chose), to exclude the last N time periods, and specify how to handle nulls.

To **Set a rolling date** for your comparison, choose **Set a rolling date**.

In the **Set a rolling date** pane that opens, choose **Relative date** and then select if you want to set the date to **Today**, **Yesterday**, or you can specify the **Filter condition** (start of or end of), **Range** (this, previous, or next), and **Period** (year, quarter, month, week, or day).

5. For **Time granularity**, choose **Day**, **Hour**, **Minute**, or **Second**.
6. (Optional) If you are filtering by using an existing parameter, instead of specific dates, choose **Use parameters**, then choose your parameter or parameters from the list. To use **Before**, **After**, or **Equals** comparisons, choose one date parameter. You can include this date in the range.

To use **Between**, enter both the start date and end date parameters separately. You can include the start date, the end date, or both in the range.

To use parameters in a filter, create them first. Usually, you create a parameter, add a control for it, and then add a filter for it. For more information, see [Parameters in Amazon Quick](#).

7. For **Null options** choose **Exclude nulls**, **Include nulls**, or **Nulls only**.

8. When finished, choose **Apply**.

Creating relative date filters in analyses

Use the following procedure to create a relative filter for a date field in an analysis.

To create a relative filter for a date field in an analysis

1. Create a new filter using a text field. For more information about creating filters, see [Adding filters](#).
2. In the **Filters** pane, choose the new filter to expand it.
3. For **Filter type**, choose **Relative dates**.
4. For **Time granularity**, choose a granularity of time that you want to filter by (days, hours, minutes).
5. For **Period**, choose a unit of time (years, quarters, quarters, months, weeks, days).
6. For **Range**, choose how you want the filter to relate to the time frame. For example, if you choose to report on months, your options are previous month, this month, month to date, last N months, and next N months.

If you choose Last N or Next N years, quarters, months, weeks, or days, enter a number for **Number of**. For example, last 3 years, next 5 quarters, last 5 days.

7. For **Null options** choose **Exclude nulls**, **Include nulls**, or **Nulls only**.
8. For **Set dates relative to**, choose one of the following options:
 - **Current date time** – If you choose this option, you can set it to **Exclude last**, and then specify the number and type of time periods.
 - **Date and time from a parameter** – If you choose this option, you can select an existing datetime parameter.
9. (Optional) If you are filtering by using an existing parameter, instead of specific dates, enable **Use parameters**, then choose your parameter or parameters from the list.

To use parameters in a filter, create them first. Usually, you create a parameter, add a control for it, and then add a filter for it. For more information, see [Parameters in Amazon Quick](#).

10. When finished, choose **Apply**.

Creating top and bottom date filters in analyses

Use the following procedure to create a top and bottom filter for a date field in an analysis.

To create a top and bottom filter for a date field in an analysis

1. Create a new filter using a text field. For more information about creating filters, see [Adding filters](#).
2. In the **Filters** pane, choose the new filter to expand it.
3. For **Filter type**, choose **Top and bottom**.
4. Select **Top** or **Bottom**.
5. For **Show**, enter the number of top or bottom items you want to show and choose a unit of time (years, quarters, months, weeks days, hours, minutes).
6. For **By**, choose a field to base the ranking on.
7. (Optional) Add another field as a tie breaker, if the field for **By** has duplicates. Choose **Tie breaker**, and then choose another field. To remove a tie breaker, use the delete icon.
8. (Optional) If you are filtering by using an existing parameter, instead of specific dates, select **Use parameters**, then choose your parameter or parameters from the list.

To use a parameter for **Top and bottom**, choose an integer parameter for the number of top or bottom items to show.

To use parameters in a filter, create them first. Usually, you create a parameter, add a control for it, and then add a filter for it. For more information, see [Parameters in Amazon Quick](#).

9. When finished, choose **Apply**.

Adding filter conditions (group filters) with AND and OR operators

In analyses, when you add multiple filters to a visual, Quick uses the AND operator to combine them. You can also add filter conditions to a single filter with the OR operator. This is called a compound filter, or filter group.

To add multiple filters using the OR operator, create a filter group. Filter grouping is available for all types of filters in analyses.

When you filter on multiple measures (green fields marked with #), you can apply the filter conditions to an aggregate of that field. Filters in a group can contain either aggregated or nonaggregated fields, but not both.

To create a filter group

1. Create a new filter in an analysis. For more information about creating filters, see [Adding filters](#).
2. In the **Filters** pane, choose the new filter to expand it.
3. In the expanded filter, choose **Add filter condition** at bottom, and then choose a field to filter on.
4. Choose the conditions to filter on.

The data type of the field that you selected determines the options available here. For example, if you chose a numeric field, you can specify the aggregation, filter condition, and values. If you chose a text field, you can choose the filter type, filter condition, and values. And if you chose a date field, you can specify the filter type, condition, and time granularity. For more information about these options, see [Filter types in Amazon Quick](#).

5. (Optional) You can add additional filter conditions to the filter group by choosing **Add filter condition** again at bottom.
6. (Optional) To remove a filter from the filter group, choose the trash-can icon near the field name.
7. When finished, choose **Apply**.

The filters appear as a group in the **Filters** pane.

Creating cascading filters

The idea behind cascading any action, such as a filter, is that choices in the higher levels of a hierarchy affect the lower levels of a hierarchy. The term *cascading* comes from the way that a cascade waterfall flows from one tier to the next.

To set up cascading filters, you need a trigger point where the filter is activated, and target points where the filter is applied. In Quick, the trigger and target points are included in visuals.

To create a cascading filter, you set up an action, not a filter. This approach is because you need to define how the cascading filter is activated, which fields are involved, and which visuals are filtered when someone activates it. For more information, including step-by-step instructions, see [Using custom actions for filtering and navigating](#).

There are two other ways to activate a filter across multiple visuals:

- **For a filter that is activated from a widget on a dashboard** – The widget is called a *sheet control*, which is a custom menu that you can add to the top of your analysis or dashboard. The most common sheet control is a drop-down list, which displays a list of options to choose from when you open it. To add one of these to your analysis, create a parameter, add a control to the parameter, and then add a filter that uses the parameter. For more information, see [Setting up parameters in Amazon Quick](#), [Using a control with a parameter in Amazon Quick](#), and [Adding filter controls to analysis sheets](#).
- **For a filter that always applies to multiple visuals** – This is a regular filter, except that you set its scope to apply to multiple (or all) visuals. This type of filter doesn't really cascade, because there is no trigger point. It always filters all the visuals that it's configured to filter. To add this type of filter to your analysis, create or edit a filter and then choose its scope: **Single visual**, **Single sheet**, or **Cross sheets**. Note the option to **Apply cross-datasets**. If this box is checked, then the filter will be applied to all visuals from different datasets that are applicable on all sheets in the filter scope. For more information, see [Filters](#).

Adding filter controls to analysis sheets

When you're designing an analysis, you can add a filter to the analysis sheet near the visuals that you want to filter. It appears in the sheet as a control that dashboard viewers can use when you publish the analysis as a dashboard. The control uses the analysis theme settings so it looks like it's part of the sheet.

Filter controls share some settings with their filters. They apply to one, some, or all of the objects on the same sheet.

Use the following sections add and customize filter controls to an analysis. To learn how to add cross-sheet controls, see [Controls](#).

Topics

- [Adding filter controls](#)
- [Pinning filter controls to the top of a sheet](#)
- [Customizing filter controls](#)
- [Sorting filter control values](#)
- [Cascading filter controls](#)

Adding filter controls

Use the following procedure to add a filter control.

To add a filter control

1. Open the [Quick console](#).
2. From the Quick homepage, choose **Analyses**, and then choose the analysis that you want to work with.
3. In the analysis, choose **Filter**.
4. If you don't already have some filters available, create one. For more information about creating filters, see [Adding filters](#).
5. In the **Filters** pane, choose the three dots to the right of the filter that you want to add a control for, and choose **Add to sheet**.

The filter control is added to the sheet, usually at the bottom. You can resize it or drag it to different positions on the sheet. You can also customize how it appears and how dashboard viewers can interact with it. For more information about customizing filter controls, see the following sections.

Pinning filter controls to the top of a sheet

Use the following procedure to pin filter controls to the top of a sheet.

To pin a control to the top of a sheet

1. On the filter control that you want to move, choose the three dots next to the pencil icon and choose **Pin to top**.

The filter is pinned to the top of the sheet and is collapsed. You can click it to expand it.

2. (Optional) To unpin the control, expand it and hover over it at the top of the sheet until three dots appear. Choose the three dots, and then choose **Move to sheet**.

Customizing filter controls

Depending on the data type of the field and the type of filter, filter controls have different settings available. You can customize how they appear in the sheet and how dashboard viewers can interact with them.

To customize a filter control

1. Choose the filter control in the sheet.
2. On the filter control, choose the pencil icon.

If the filter control is pinned to the top of the sheet, expand it and hover your cursor over it until the three dots appear. Choose the three dots, and then choose **Edit**.

3. In the **Format control** pane that opens, do the following:
 - a. For **Display name**, enter a name for the filter control.
 - b. (Optional) To hide the display name from the filter control, clear the check box for **Show title**.
 - c. For **Title font size**, choose the title font size that you want to use. The options range from extra small to extra large. The default setting is medium.

The remaining steps depend on the type of field the control is referencing. For options by filter type, see the following sections.

Date filters

If your filter control is from a date filter, use the following procedure to customize the remaining options.

To customize further options for a date filter

1. In the **Format control** pane, for **Style**, choose one of the following options:
 - **Date picker – range** – Displays a set of two fields to define a time range. You can enter a date or time, or you can choose a date from the calendar control. You can also customize how you want the dates to appear in the control by entering a date token for **Date format**. For more information, see [Customizing date formats in Quick](#).
 - **Date picker – relative** – Displays settings like the time period, its relation to the current date and time, and the option to exclude time periods. You can also customize how you want the dates to appear in the control by entering a date token for **Date format**. For more information, see [Customizing date formats in Quick](#).
 - **Text field** – Displays a box where you can enter the top or bottom *N* date.

Helper text is included in the text field control by default, but you can choose to remove it by clearing the **Show helper text in control** option.

By default, Quick visuals are reloaded whenever a change is made to a control. For Calendar and Relative date picker controls, authors can add an **Apply** button to a control that delays visual reload until the user chooses **Apply**. This allows users to make multiple changes at a time without additional queries. This setting can be configured with the **Show an apply button** checkbox in the **Control options** section of the **Format control** pane.

2. When finished, choose **Apply**.

Text filters

If your filter control is from a text filter, for example dimensions, categories, or labels, use the following procedure to customize the remaining options.

To customize further options for a text filter

1. In the **Format control** pane, for **Style**, choose one of the following options:
 - **Dropdown** – Displays a dropdown list with buttons that you can use to select a single value.

When you select this option, you can choose the following options for **Values**:

- **Filter** – Displays all the values that are available in the filter.
- **Specific values** – Enables you to enter the values to display, one entry per line.

You can also choose to **Hide Select all option from the control values**. This removes the option to select or clear the selection of all values in the filter control.

- **Dropdown - multiselect** – Displays a dropdown list with boxes that you can use to select multiple values.

When you select this option, you can choose the following options for **Values**:

- **Filter** – Displays all the values that are available in the filter.
- **Specific values** – Enables you to enter the values to display, one entry per line.

By default, Quick visuals are reloaded whenever a change is made to a control. For Multiselect dropdown controls, authors can add an **Apply** button to a control that delays visual reload until the user chooses **Apply**. This allows users to make multiple changes at

a time without additional queries. This setting can be configured with the **Show an apply button** checkbox in the **Control options** section of the **Format control** pane.

- **List** – Displays a list with buttons that you can use to select a single value.

When you select this option, you can choose the following options for **Values**:

- **Filter** – Displays all the values that are available in the filter.
- **Specific values** – Enables you to enter the values to display, one entry per line.

You can also choose the following:

- **Hide search bar when control is on sheet** – Hides the search bar in the filter control, so users can't search for specific values.
- **Hide Select all option from the control values** – Removes the option to select or clear the selection of all values in the filter control.
- **List - multiselect** – Displays a list with boxes that you can use to select multiple values.

When you select this option, you can choose the following options for **Values**:

- **Filter** – Displays all the values that are available in the filter.
- **Specific values** – Enables you to enter the values to display, one entry per line.

You can also choose the following:

- **Hide search bar when control is on sheet** – Hides the search bar in the filter control, so users can't search for specific values.
- **Hide Select all option from the control values** – Removes the option to select or clear the selection of all values in the filter control.
- **Text field** – Displays a text box where you can enter a single entry. Text fields support up to 79950 characters.

When you select this option, you can choose the following:

- **Show helper text in control** – Removes the helper text in text fields.
- **Text field - multiline** – Displays a text box where you can enter multiple entries. Multiline text fields support up to 79950 characters across all entries.

When you select this option, you can choose the following:

- For **Separate values by**, choose how you want to separate values you enter into the filter control. You can choose to separate values by a line break, comma, pipe (|), or semicolon.

- **Show helper text in control** – Removes the helper text in text fields.

2. When finished, choose **Apply**.

Numeric filters

If your filter control is from a numeric filter, use the following procedure to customize the remaining options.

To customize further options for a numeric filter

1. In the **Format control** pane, for **Style**, choose one of the following options:

- **Dropdown** – Displays a list where you can select a single value.

When you select this option, you can choose the following options for **Values**:

- **Filter** – Displays all the values that are available in the filter.
- **Specific values** – Enables you to enter the values to display, one entry per line.

You can also choose to **Hide Select all option from the control values**. This removes the option to select or clear the selection of all values in the filter control.

- **Filter** – Displays all the values that are available in the filter.
- **Specific values** – Enables you to enter the values to display, one entry per line.
- **Hide Select all option from the control values** – Removes the option to select or clear the selection of all values in the filter control.
- **List** – Displays a list with buttons that enable selecting a single value.

When you select this option, you can choose the following options for **Values**:

- **Filter** – Displays all the values that are available in the filter.
- **Specific values** – Enables you to enter the values to display, one entry per line.

You can also choose the following:

- **Hide search bar when control is on sheet** – Hides the search bar in the filter control, so users can't search for specific values.
- **Hide Select all option from the control values** – Removes the option to select or clear the selection of all values in the filter control.

- **Slider** – Displays a horizontal bar with a toggle that you can slide to change the value. If you have a ranged filter for values between a minimum and a maximum, the slider provides a toggle for each number. For sliders, you can specify the following options:
 - **Minimum value** – Displays the smaller value at the left of the slider.
 - **Maximum value** – Displays the larger value at the right of the slider.
 - **Step size** – Enables you to set the number of increments that the bar is divided into.
- **Text box** – Displays a box where you can enter the value. When you select this option, you can choose the following:
 - **Show helper text in control** – Removes the helper text in text fields.

2. When finished, choose **Apply**.

Sorting filter control values

Custom sort gives authors control over how values appear in dropdown and list filter controls. By default, control values are sorted alphabetically in ascending order. With custom sort, you can display values in a specific business order (such as fiscal quarters or priority levels) or sort by related metrics (such as sorting regions by total sales).

Custom sort is available for **Dropdown** (single select and multiselect) and **List** (single select and multiselect) control styles. The available sort options depend on whether the control uses specific values or values from a dataset column.

Note

Custom sort is not available for date type columns at this time. To sort date values in a logical order, use the sort by another field option with a date-related field.

When to use custom sort

Use custom sort when you need to do the following:

- **Display values in business logic order** – For example, priority levels (Critical, High, Medium, Low), fiscal quarters (Q1, Q2, Q3, Q4), or custom status workflows (New, In Progress, Review, Complete).
- **Sort by related metrics** – For example, products sorted by total sales, regions sorted by customer count, categories sorted by average revenue, or departments sorted by headcount.

Configuring custom sort

The sort configuration options depend on whether the control uses specific values or values from a dataset column.

Controls with specific values

When a control uses specific values that you entered manually, the following sort options are available in the **Sort** section of the **Format control** pane.

To configure sort for a specific values control

1. Choose the filter control in the sheet.
2. Choose the pencil icon to open the **Format control** pane.
3. In the **Format control** pane, locate the **Sort** section.
4. Choose one of the following options:
 - **Ascending (A–Z, 0–9)** – Sorts values in ascending order. This is the default.
 - **Descending (Z–A, 9–0)** – Sorts values in descending order.
 - **User-defined order** – Displays the values in the exact order you entered them. This preserves your custom ordering without any automatic sorting.

Note

When a control has values from both specific values and a source entity (such as a filter or parameter), the combined list is sorted together. For user-defined order, values are appended in the order they were entered. If a value from the source cannot be sorted based on the current configuration, it is appended at the end of the list.

Controls with values from a dataset column

When a control displays values from a dataset column, you can sort by that field or another field in the dataset using an aggregation function. This is useful when you want to order control values by a related metric, such as sorting a list of products by total sales.

To configure sort for a column-based control

1. Choose the filter control in the sheet.

2. Choose the pencil icon to open the **Format control** pane.
3. In the **Format control** pane, locate the **Sort** section.
4. Configure the following options:
 - **Sort by control column** – Sorts the control values based on the column that the control is tied to.
 - For **Sort direction**, choose **Ascending** or **Descending**.
 - For **Aggregation**, choose an aggregation function or choose **No aggregation** to sort by the raw column values.
 - **Sort by another field** – Sorts the control values based on a different column in the dataset.
 - For **Sort by field**, choose a column from the dataset. For calculated fields, only scalar (non-aggregated) calculated columns are available.
 - For **Aggregation**, choose an aggregation function to apply to the sort field. The available functions depend on the data type of the selected column. For numeric fields, you can use Sum, Average, Count, Distinct count, Min, Max, Median, Percentile, Var, Stdev, and other statistical functions. For non-numeric fields, you can use Count and Distinct count.
 - For **Sort direction**, choose **Ascending** or **Descending**.

Cross-sheet filter controls

You can configure sort order for cross-sheet filter controls through the cross-sheet settings. The sort configuration applies to all instances of the control across sheets.

To configure sort for a cross-sheet filter control

1. Open the cross-sheet settings for the filter.
2. In the default control configuration section, locate the **Sort** settings.
3. Configure the sort options as described in the preceding sections.
4. Choose **Apply**.

Note

Dashboard controls inherit the sort configuration from the analysis. Sort configuration changes are not available to dashboard readers.

Cascading filter controls

You can limit the values displayed in the control, so they only show values that are valid for what is selected in other controls. This is called a cascading control.

When creating cascading controls, the following limitations apply:

1. Cascading controls must be tied to dataset columns from the same dataset.
2. The child control must be a dropdown or list control.
3. For parameter controls, the child control must be linked to a dataset column.
4. For filter controls, the child control must be linked to a filter (instead of showing only specific values).
5. The parent control must be one of the following:
 - a. A string, integer, or numeric parameter control. DateTime parameter control is not supported.
 - b. A string filter control (excluding top-bottom filters).
 - c. A non-aggregated numeric filter control.
 - d. A date filter control (excluding top-bottom filters).

To create a cascading control

1. Choose **Show relevant values only**. Note that this option might not be available for all filter control types.
2. In the **Show relevant values only** pane that opens, choose one or more controls from the available list.
3. Choose a field to match the value to.
4. Choose **Update**.

Editing filters

You can edit filters at any time in a dataset or analysis.

You can't change the field a filter applies to. To apply a filter to a different field, create a new filter instead.

Use the following procedures to learn how to edit filters.

Editing filters in datasets

Use the following procedure to edit filters in datasets.

To edit a filter in a dataset

1. Open the [Quick console](#).
2. From the Quick homepage, choose **Data** at left.
3. Under the **Datasets** tab, choose the dataset that you want, and then choose **Edit dataset**.
4. On the data preparation page that opens, choose **Filters** at lower left.
5. Choose the filter that you want to edit.
6. When finished editing, choose **Apply**.

Editing filters in analyses

Use the following procedure to edit filters in analyses.

To edit a filter in an analysis

1. Open the [Quick console](#).
2. From the Quick homepage, choose **Analyses** at left.
3. On the **Analyses** page, choose the analysis that you want to work with.
4. In the analysis, choose the **Filter** icon shown to open the **Filters** pane.
5. Choose the filter that you want to edit.
6. When finished editing, choose **Apply**.

Enabling or disabling filters

You can use the filter menu to enable or disable a filter in a dataset or an analysis. When you create a filter, it's enabled by default. Disabling a filter removes the filter from the field, but it doesn't delete the filter from the dataset or analysis. Disabled filters are grayed out in the filters pane. If you want to re-apply the filter to the field, you can simply enable it.

Use the following procedures to learn how to enable or disable filters.

Disabling filters in datasets

Use the following procedure to disable filters in datasets.

To disable a filter in a dataset

1. From the Quick homepage, choose **Data** at left.
2. Under the **Datasets** tab, choose the dataset that you want, and then choose **Edit dataset**.
3. On the data preparation page that opens, choose **Filters** at lower left.
4. In the **Filters** pane at left, choose the three dots to the right of the filter that you want to disable, and then choose **Disable**. To enable a filter that was disabled, choose **Enable**.

Disabling filters in analyses

Use the following procedure to disable filters in analyses.

To disable a filter in an analysis

1. Open the [Quick console](#).
2. From the Quick homepage, choose **Analyses**.
3. On the **Analyses** page, choose the analysis that you want to work with.
4. In the analysis, choose the **Filter** icon to open the **Filters** pane.
5. In the **Filters** pane that opens, choose the three dots to the right of the filter that you want to disable, and then choose **Disable**. To enable a filter that was disabled, choose **Enable**.

Deleting filters

You can delete filters at any time in a dataset or analysis. Use the following procedures to learn how.

Deleting filters in datasets

Use the following procedure to delete filters in datasets.

To delete a filter in a dataset

1. Open the [Quick console](#).
2. From the Quick homepage, choose **Data**.
3. Under the **Datasets** tab, choose the dataset that you want, and then choose **Edit dataset**.
4. On the data preparation page that opens, choose **Filters** at lower left.
5. Choose the filter that you want to delete, and then choose **Delete filter**.

Deleting filters in analyses

Use the following procedure to delete filters in analyses.

To delete a filter in an analysis

1. Open the [Quick console](#).
2. From the Quick homepage, choose **Analyses**.
3. On the **Analyses** page, choose the analysis that you want to work with.
4. In the analysis, choose the **Filter** icon to open the **Filters** pane.
5. Choose the filter that you want to delete, and then choose **Delete filter**.

Previewing tables in a dataset

You can preview each individual data table within a dataset. When you choose a data table to preview, a read-only preview of the table appears in a new tab in the data preview section. You can have multiple table preview tabs open at once.

You can only preview tables that you have access to in a dataset. If a table doesn't appear in the top half of the data preparation space, you can't preview the table.

The **Dataset** tab contains all transformations, like new columns or filters. Table preview tabs don't show any of your transforms.

To preview a data table

1. On the Quick homepage, choose **Data** at left.
2. In the **Data** tab, choose the dataset that you want, and choose **Edit dataset**.
3. Choose the data table that you want to preview, choose the down arrow to open the menu, and choose **Show table preview**.

Using SQL to customize data

When you create a dataset or prepare your data for use in an analysis, you can customize the data in the query editor.

The query editor is made up of multiple components, as follows:

- **Query mode** – At the top left, you can choose between direct query or SPICE query modes:

- **Direct query** – To run the SELECT statement directly against the database
- **SPICE** – To run the SELECT statement against data that was previously stored in memory
- **Fields** – Use this section to disable fields you want to remove from the final dataset. You can add calculated fields in this section, and augment your data with SageMaker AI
- **Query archive** – Use this section to find previous version of your SQL queries.
- **Filters** – Use this section to add, edit, or remove filters.
- **Schema explorer** – This section only appears while you are editing SQL. You can use it to explore your schemas, tables, fields, and data types.
- **SQL editor** – Use this to edit your SQL. The SQL editor, which offers syntax highlighting, basic autocomplete, autoindent, and line numbering. You can specify a SQL query only for datasets that come from data sources compatible with SQL. Your SQL must conform to the target database requirements regarding syntax, capitalization, command termination, and so on. If you prefer, you can instead paste SQL from another editor.
- **Data workspace** – When the SQL editor is closed, the data workspace displays at top right with a grid background. Here you can see a graphical representation of your data objects, including queries, tables, files, and joins created in the join editor.

To view details about each table, use the data source options menu and choose **Table details** or **Edit SQL Query**. Details display for table name and alias, schema, data source name, and data source type. For upload settings on a file, choose **Configure upload settings** from the data source options menu to view or change the following settings:

- **Format** – the file format, CSV, CUSTOM, CLF, and so on
- **The starting row** – the row to start with
- **The text qualifier** – double quote or single quote
- **Header** – indicates if the file includes a header row
- **Preview rows** – A preview of the sampled rows appear at bottom right when the join configuration editor isn't in use.
- **Join configuration editor** – The join editor opens when you have more than one data object in the data workspace. To edit a join, you select the join icon between two tables (or files). Choose a join type and the fields to join on, by using the join configuration panel at the bottom of the screen. Then choose **Apply** to create the join. You must complete all joins before you can save your work.

To add more queries, tables, or files, use the **Add data** option above the workspace.

Creating a basic SQL query

Use the following procedure to connect to a data source by using a custom SQL query.

To create a basic SQL query

1. Create a new data source and validate the connection.
2. Fill in the options necessary to connection, however you don't need to select a schema or a table.
3. Choose **Use custom SQL**.
4. (Optional) You can enter your query in the SQL editor, or continue on to the next step to use the full-screen version. To enter it now, create a name for the query. Then type or paste a SQL query into the editor. The SQL editor offers syntax highlighting, basic autocomplete, autoindent, and line numbering.

(Optional) Choose **Confirm query** to validate it and view settings for direct query, SPICE memory, and SageMaker AI settings.

5. Choose **Edit/Preview data**. The full query editor appears with the SQL editor displayed. The query is processed and a sample of the query results displays in the data preview pane. You can make changes to the SQL and confirm them by choosing **Apply**. When you are done with the SQL, choose **Close** to continue.
6. At the top, enter a name for the dataset. Then choose **Save & visualize**.

Modifying existing queries

To update a SQL query

1. Open the dataset that you want to work with.
2. In the workspace with the grid, locate the box-shaped object that represents the existing query.
3. Open the options menu on the query object and choose **Edit SQL query**. If this option doesn't appear in the list, the query object isn't based on SQL.

To view previous versions of queries, open the **Query archive** at left.

Adding geospatial data

You can flag geographic fields in your data, so that Amazon Quick Sight can display them on a map. Amazon Quick Sight can chart latitude and longitude coordinates. It also recognizes geographic components such as country, state or region, county or district, city, and ZIP code or postal code. You can also create geographic hierarchies that can disambiguate similar entities, for example the same city name in two states.

Note

Geospatial charts in Amazon Quick Sight aren't currently supported in some AWS Regions, including in China. We are working on adding support for more Regions.

Use the following procedures to add geospatial data types and hierarchies to your dataset.

To add geospatial data types and hierarchies to your dataset

1. On the data preparation page, label the geographic components with the correct data type.

There are several ways to do this. One is to choose the field under **Fields** and use the ellipses icon (...) to open the context menu.

Then choose the correct geospatial data type.

You can also change the data type in the work area with the data sample. To do this, choose the data type listed under the field name. Then choose the data type you want to assign.

2. Verify that all geospatial fields necessary for mapping are labeled as geospatial data types. You can check this by looking for the place marker icon. This icon appears under the field names across the top of the page, and also in the **Fields** pane on the left.

Also check the name of the data type, for example latitude or country.

3. (Optional) You can set up a hierarchy or grouping for geographical components (state, city), or for latitude and longitude coordinates. For coordinates, you must add both latitude and longitude to the geospatial field wells.

To create a hierarchy or grouping, first choose one of these fields in the **Fields** pane. Each field can only belong to one hierarchy. It doesn't matter which one you choose first, or what order you add the fields in.

Choose the ellipsis icon (...) next to the field name. Then choose **Add to a hierarchy**.

4. On the **Add field to hierarchy** screen, choose one of the following:
 - Choose **Create a new geospatial hierarchy** to create a new grouping.
 - Choose **Add to existing geospatial hierarchy** to add a field to a grouping that already exists. The existing hierarchies displayed include only those of matching geospatial types.

Choose **Add** to confirm your choice.

5. On the **Create hierarchy** screen, name your hierarchy.

If you are creating a latitude and longitude grouping, the **Create hierarchy** screen appears. Depending on whether you chose latitude or longitude in the previous steps, either latitude or longitude displays on this screen. Make sure that your latitude field shows under **Field to use for latitude**. Also make sure that your longitude shows under **Field to use for longitude**.

For geographical components, the **Create hierarchy** screen has two choices:

- Choose **This hierarchy is for a single country** if your data only contains one country. Choose the specific country from the list. Your data doesn't need to contain every level of the hierarchy. You can add fields to the hierarchy in any order.
- Choose **This hierarchy is for multiple countries** if your data contains more than one country. Choose the field that contains the country names.

For either hierarchy type, choose **Update** to continue.

6. Continue by adding as many fields to the hierarchy as you need to.

Your geospatial groupings appear in the **Fields** pane.

Changing a geospatial grouping

You can change a geospatial hierarchy or grouping that exists in a dataset.

Use the following procedure to edit or disband a geospatial hierarchy.

To edit or disband a geospatial hierarchy

1. Open the dataset. In the **Fields** pane, choose the hierarchy name.

2. Choose the ellipsis icon (...), then choose one of the following options.

Choose **Disband hierarchy** to remove the hierarchy from the dataset. You can't undo this operation. However, you can recreate your hierarchy or grouping by starting again at step 1. Disbanding the hierarchy doesn't remove any fields from the dataset.

Choose **Edit hierarchy** to make changes to the hierarchy. Doing this reopens the creation screens, so you can make different choices in rebuilding your hierarchy.

Geospatial troubleshooting

Use this section to discover the Amazon Quick Sight requirements for correctly processing geospatial data. If Amazon Quick Sight doesn't recognize your geospatial data as geospatial, use this section to help troubleshoot the issue. Make sure that your data follows the guidelines listed, so that it works in geospatial visuals.

Note

Geospatial charts in Amazon Quick Sight currently aren't supported in some AWS Regions, including in China. We are working on adding support for more Regions. If your geography follows all the guidelines listed here, and still generates errors, contact the Amazon Quick Sight team from within the Amazon Quick Sight console.

Topics

- [Geocoding issues](#)
- [Issues with latitude and longitude](#)
- [Supported administrative areas and postal codes by country](#)

Geocoding issues

Amazon Quick Sight geocodes place names into latitude and longitude coordinates. It uses these coordinates to display place names on the map. Amazon Quick Sight skips any places that it can't geocode.

For this process to work properly, your data must include at least the country. Also, there can't be duplicate place names inside of a parent place name.

A few issues prevent place names from showing up on a map chart. These issues include unsupported, ambiguous, or invalid locations, as described following.

Topics

- [Issues with unsupported areas](#)
- [Issues with ambiguous locations](#)
- [Issues with invalid geospatial data](#)
- [Issues with the default country in geocoding](#)

Issues with unsupported areas

To map unsupported locations, include latitude and longitude coordinates in your data. Use these coordinates in the geospatial field well to make locations show on a map chart.

Issues with ambiguous locations

Geospatial data can't contain ambiguous locations. For example, suppose that the data contains a city named **Springfield**, but the next level in the hierarchy is country. Because multiple states have a city named **Springfield**, it isn't possible to geocode the location to a specific point on a map.

To avoid this problem, you can add enough geographical data to indicate what location should show on a map chart. For example, you can add a state level into your data and its hierarchy. Or, you might add latitude and longitude.

Issues with invalid geospatial data

Invalid geospatial data occurs when a place name (a city, for example) is listed under an incorrect parent (a state, for example). This issue might be a simple misspelling, or data entry error.

Note

Amazon Quick Sight doesn't support regions (for example, West Coast or South) as geospatial data. However, you can use a region as a filter in a visual.

Issues with the default country in geocoding

Make sure that you are using the correct default country.

The default for each hierarchy is based on the country or country field that you choose when you create the hierarchy.

To change this default, you can return to the **Create hierarchy** screen. Then edit or create a hierarchy, and choose a different country.

If you don't create a hierarchy, your default country is based on your AWS Region. For details, see the following table.

Region	Default country
US West (Oregon) Region	US
US East (Ohio) Region	
US East (N. Virginia) Region	
Asia Pacific (Singapore)	Singapore
Asia Pacific (Sydney)	Australia
Europe (Ireland) Region	Ireland

Issues with latitude and longitude

Amazon Quick Sight uses latitude and longitude coordinates in the background to find place names on a map. However, you can also use coordinates to create a map without using place names. This approach also works with unsupported place names.

Latitude and longitude values must be numeric. For example, the map point indicated by **28.5383355 -81.3792365** is compatible with Amazon Quick Sight. But **28° 32' 18.0096' ' N 81° 22' 45.2424' ' W** is not.

Topics

- [Valid ranges for latitude and longitude coordinates](#)
- [Using coordinates in degrees, minutes, and seconds \(DMS\) format](#)

Valid ranges for latitude and longitude coordinates

Amazon Quick Sight supports latitude and longitude coordinates within specific ranges.

Coordinate	Valid range
Latitude	Between -90 and 90
Longitude	Between -180 to 180

Amazon Quick Sight skips any data outside these ranges. Out-of-range points can't be mapped on a map chart.

Using coordinates in degrees, minutes, and seconds (DMS) format

You can use a calculated field with a formula to create a numeric latitude and longitude out of character strings. Use this section to find different ways that you can create calculated fields in Amazon Quick Sight, to parse GPS latitude and longitude into numeric latitude and longitude.

The following sample converts latitude and longitude to numeric format from separate fields. For example, suppose that you parse **51° 30' 26.4636' ' N 0° 7' 39.9288' ' W** using space as a delimiter. In this case, you can use something like the following sample to convert the resulting fields to numeric latitude and longitude.

In this example, the seconds are followed by two single quotation marks. If your data has a double quotation mark instead, then you can use `strlen(LatSec)-1` instead of `strlen(LatSec)-2`.

```
/*Latitude*/
  ifelse(
    LatDir = "N",
    parseInt(split(LatDeg, "°", 1)) +
      (parseFloat(split(LatMin, "'", 1) ) /60) +
      (parseFloat(substring(LatSec, 1, strlen(LatSec)-2) ) ) /3600),
    (parseInt(split(LatDeg, "°", 1)) +
      (parseFloat(split(LatMin, "'", 1) ) /60) +
      (parseFloat(substring(LatSec, 1, strlen(LatSec)-2) ) ) /3600)) * -1
  )

/*Longitude*/
  ifelse(
```

```

LongDir = "E",
parseInt(split(LongDeg, "°", 1)) +
    (parseFloat(split(LongMin, "'", 1) ) /60) +
    (parseFloat((substring(LongSec, 1, strlen(LongSec)-2) ) ) /3600),
(parseInt(split(LongDeg, "°", 1)) +
    (parseFloat(split(LongMin, "'", 1) ) /60) +
    (parseFloat((substring(LongSec, 1, strlen(LongSec)-2) ) ) /3600)) * -1
)

```

If your data doesn't include the symbols for degree, minute and second, the formula looks like the following.

```

/*Latitude*/
ifelse(
    LatDir = "N",
    (LatDeg + (LatMin / 60) + (LatSec / 3600)),
    (LatDeg + (LatMin / 60) + (LatSec / 3600)) * -1
)

/*Longitude*/
ifelse(
    LongDir = "E",
    (LongDeg + (LongMin / 60) + (LongSec / 3600)),
    (LongDeg + (LongMin / 60) + (LongSec / 3600)) * -1
)

```

The following sample converts **53°21'N 06°15'W** to numeric format. However, without the seconds, this location doesn't map as accurately.

```

/*Latitude*/
ifelse(
    right(Latitude, 1) = "N",
    (parseInt(split(Latitude, '°', 1)) +
        parseFloat(substring(Latitude, (locate(Latitude, '°',3)+1), 2) ) / 60) ,
    (parseInt(split(Latitude, '°', 1)) +
        parseFloat(substring(Latitude, (locate(Latitude, '°',3)+1), 2) ) / 60) * -1
)

/*Longitude*/
ifelse(
    right(Longitude, 1) = "E",

```

```
(parseInt(split(Longitude, '°', 1)) +
  parseDecimal(substring(Longitude, (locate(Longitude, '°',3)+1), 2) ) / 60) ,
(parseInt(split(Longitude, '°', 1)) +
  parseDecimal(substring(Longitude, (locate(Longitude, '°',3)+1), 2) ) / 60) *
-1
)
```

The formats of GPS latitude and longitude can vary, so customize your formulas to match your data. For more information, see the following:

- [Degrees Minutes Seconds to Decimal Degrees](#) on LatLong.net
- [Converting Degrees/Minutes/Seconds to Decimals using SQL](#) on Stack Overflow
- [Geographic Coordinate Conversion](#) on Wikipedia

Supported administrative areas and postal codes by country

The following is a list of supported administrative areas by country.

Supported administrative areas

Country name	Country code	Country	State	County	City
Aruba	ABW	Country	Regions	Zones	
Afghanistan	AFG	Country	Wilayat	Wuleswali	Localities/Urban Areas
Angola	AGO	Country	Provinces /Províncias	Municipios	Localities/Urban Areas
Anguilla	AIA	Country	Parishes		
Albania	ALB	Country	Qarqe/Qark	Communes: Bashki	Njësi/Loc

Country name	Country code	Country	State	County	City
					alities/ Urban Areas
Andorra	AND	Country	Parishes/ Parròquies	Localities/ Urban Areas	
United Arab Emirates	ARE	Country	Emirates	Municipalities	Cities/Localities/ Urban Areas
Argentina	ARG	Country	Provinces/ Provincias	Departamentos/ Departments	Comunas/ Barrios
Armenia	ARM	Country	Provinces/ Marzpet		Localities/ Urban Areas
American Samoa	ASM	Country	Districts	Counties	Villages
Antarctica	ATA	Country			
French Southern Territories	ATF	Country	Districts		
Antigua and Barbuda	ATG	Country	Parishes		Localities/ Urban Areas

Country name	Country code	Country	State	County	City
Australia	AUS	Country	States	Local Government Areas	Suburbs/ Urban Centers
Austria	AUT	Country	States/ Bundesländer	Districts/ Bezirke	Municipalities/ Gemeinden/ Urban Areas/ Stadtteil
Azerbaijan	AZE	Country	Regions/ İqtisadi Rayonlar	Districts/ Rayonlar	Localities/ Urban Areas
Burundi	BDI	Country	Provinces	Communes	Localities/ Urban Areas
Belgium	BEL	Country	Regions/ Gewest	Provinces/ Provincie	Districts/ Arrondissements/ Municipalities/ Communes
Benin	BEN	Country	Departments	Communes	Localities/ Urban Areas

Country name	Country code	Country	State	County	City
Bonaire, Sint Eustasius, and Saba	BES	Country	Municipalities		Localities/Urban Areas
Burkina Faso	BFA	Country	Regions	Provinces	Communes/Localities/Urban Areas
Bangladesh	BGD	Country	Divisions/Bibhag	Districts/Zila	Subdistricts/Upzila/Localities/Urban Areas
Bulgaria	BGR	Country	Oblasts	Obshtina	Localities/Urban Areas
Bahrain	BHR	Country	Governorates	Constituencies	Localities
Bahamas	BHS	Country	Island Groups	Districts	Towns
Bosnia and Herzegovina	BIH	Country	Federation/Republika	Kanton	Opština/Localities/Urban Areas

Country name	Country code	Country	State	County	City
Saint Barthélemy	BLM	Country			Localities/Urban Areas
Belarus	BLR	Country	Voblasts	Rayon	Selsoviet /Localities/ Urban Areas
Belize	BLZ	Country	Districts	Constituencies	Localities/Urban Areas
Bermuda	BMU	Country	Parishes		Localities/Urban Areas
Bolivia	BOL	Country	Provinces /Provincias	Departamentos/ Departments	Municipalities/ Municipios/Localities /Urban Areas
Brazil	BRA	Country	Provinces /States/ Unidades	Municipalities/ Municipios	Localities/Urban Areas

Country name	Country code	Country	State	County	City
Barbados	BRB	Country	Parishes		Localities/ Urban Areas
Brunei	BRN	Country	Districts/ Dawair	Subdistricts/ Mukim	Villages/ Kampung/ Localities/ Urban Areas
Bhutan	BTN	Country	Districts/ Dzongkhag		Localities/ Urban Areas
Bouvet Island	BVT	Country			
Botswana	BWA	Country	Districts	Subdistricts	Localities/ Urban Areas
Central African Republic	CAF	Country	Regions	Prefectures	Sub Prefectures/ Communes

Country name	Country code	Country	State	County	City
Canada	CAN	Country	Provinces/ Territories	Census Divisions	Census Subdivisions/ Localities/ Urban Areas
Switzerland	CHE	Country	Cantons/ Kanton/ Cantone/ Chantun	District/ Bezirk/ Distretto/ Circuli	"Commune/ Gemeinde/ Comune/ Cumün/ Localities/ Urban Areas"
Chile	CHL	Country	Regions/ Regiones	Province/ Provincias	Communes/ Comunas/ Localities/ Urban Areas
China, People's Republic of	CHN	Country	Provinces	Prefectures	Cities/ Counties
Cote d'Ivoire	CIV	Country	Districts	Regions	Departments/ Sub Prefectures

Country name	Country code	Country	State	County	City
Cameroon	CMR	Country	Provinces /Regions	Departments	Arrondissements/ Cities
Congo, Democratic Republic of	COD	Country	Provinces	Districts	Localities/ Urban Areas
Congo, Republic of	COG	Country	Departments		Communes/ Arrondissements
Cook Islands	COK	Country	Island Councils		
Colombia	COL	Country	Departments	Municipios	Localities/ Urban Areas
Comoros	COM	Country	Autonomous Islands/ îles Autonomes		Villes/Villages
Clipperton Island	CPT	Country			
Cape Verde	CPV	Country	Ilhas	Concelhos	Localities/ Urban Areas

Country name	Country code	Country	State	County	City
Costa Rica	CRI	Country	Provincias	Cantons	Distritos/ Localities/ Urban Areas
Cuba	CUB	Country	Provincias	Municipios	Localities/ Urban Areas
Curaçao	CUW	Country			Localities/ Urban Areas
Cayman Islands	CYM	Country	Districts		
Cyprus	CYP	Country	Districts/ Eparchies	Municipalities/ Dimos	Localities/ Urban Areas/ Sinikia
Czech Republic	CZE	Country	Regions/ Kraj	Municipalities/ Orp	Obec/ Mesto
Germany	DEU	Country	Bundesland/ States	Kreis/ Districts	Gemeinde/ Municipalities/ Stadtteil/ Localities/ Urban Areas

Country name	Country code	Country	State	County	City
Djibouti	DJI	Country	Regions		Localities/Urban Areas
Dominica	DMA	Country	Parishes		Localities/Urban Areas
Denmark	DNK	Country	Regions	Provinces	Municipalities/Localities/Urban Areas
Dominican Republic	DOM	Country	Regions/ Regiones	Provinces/ Provincias	Municipalities/ Municipios/Localities/Urban Areas
Algeria	DZA	Country	Provinces/ Wilayas	Districts	Municipalities/Baladiyas/Localities/Urban Areas

Country name	Country code	Country	State	County	City
Ecuador	ECU	Country	Provinces	Cantons	Parishes/ Localities/ Urban Areas
Egypt	EGY	Country	Governorates/ Muhafazat	Municipal Divisions /Markaz	Towns/ Cities/ Sub Municipal Divisions
Eritrea	ERI	Country	Regions/ Zoba	Districts/ Subzobas	Localities/ Urban Areas
Spain	ESP	Country	Autonomous Communities/ Comunidades Autónomas	Provincias	Municipios/ Localities/ Urban Areas
Estonia	EST	Country	Maakond	Omavalitsused/ Linn/ Vald	Küla/ Localities/ Urban Areas
Ethiopia	ETH	Country	Regions/ Kililoch	Zones/ Zonouch	Localities/ Urban Areas

Country name	Country code	Country	State	County	City
Finland	FIN	Country	Regions/ M aakunta	Sub- Regions/ Seutuk unta	Municipal ities/Kun ta/Locali ties/ Urban Areas
Fiji	FJI	Country	Divisions	Provinces	Districts/ Villages
Falkland Islands	FLK	Country			
France	FRA	Country	Regions	Départeme nts	Arrondiss ements/ Cantons
Faroe Islands	FRO	Country	Regions/ Syslur	Municipal ities/ Kom munur	Localitie s/Urban Areas
Federated States of Micronesia	FSM	Country	States		
Gabon	GAB	Country	Provinces	Departmer ts	Localitie s/Urban Areas

Country name	Country code	Country	State	County	City
United Kingdom	GBR	Country	Nations	Counties	Districts /Localities/ Urban Areas
Georgia	GEO	Country	Regions/ Mkhare	Municipalities/ Municipality	Localities/ Urban Areas
Ghana	GHA	Country	Regions	Districts	Localities/ Urban Areas
Gibraltar	GIB	Country			Localities/ Urban Areas
Guinea	GIN	Country	Regions	Prefectures	Sub Prefectures/ Localities/ Urban Areas
Guadeloupe	GLP	Country	Arrondissements	Communes	Localities/ Urban Areas

Country name	Country code	Country	State	County	City
Gambia	GMB	Country	Regions	Districts	Localities/Urban Areas
Guinea Bissau	GNB	Country	Regions	Sectors	Localities/Urban Areas
Equatorial Guinea	GNQ	Country	Regions	Provincias	Distritos/Localities/Urban Areas
Greece	GRC	Country	Regions/Peripheries	Regional Units/Peri Enotities	Municipalities/Domoi/Localities/Urban Areas
Grenada	GRD	Country	States	Parishes/Dependencies	Localities/Urban Areas
Greenland	GRL	Country	Municipalities/Kommunia		

Country name	Country code	Country	State	County	City
Guatemala	GTM	Country	Departments/ Departamentos	Municipalities/ Municipios	Localities/ Urban Areas
French Guiana	GUF	Country	Arrondissements	Communes	Localities/ Urban Areas
Guam		Country = USA	States	Districts	
Guyana	GUY	Country	Regions	Neighborhood Councils	Localities/ Urban Areas
Hong Kong	HKG	Country	Districts	Subdistricts	Localities/ Urban Areas
Heard and McDonald Islands	HMD	Country			
Honduras	HND	Country	Departments/ Departamentos	Municipalities/ Municipios	Localities/ Urban Areas
Croatia	HRV	Country	Counties	Municipalities	Localities/ Urban Areas

Country name	Country code	Country	State	County	City
Haiti	HTI	Country	Departments/ Départements	Districts/ Arrondissements	Communes/ Localities/ Urban Areas
Hungary	HUN	Country	Regiok	Megyék	Járások/ Városok
Indonesia	IDN	Country	Provinces/ Provinsi	Regency/ Kabupaten	Districts/ Kecamatan/ Localities/ Urban Areas
India	IND	Country	States/ Territories	Districts	Subdistricts/ Towns/ Localities/ Urban Areas
British Indian Ocean Territory	IOT	Country			
Ireland	IRL	Country	Regions	Counties	Electoral Divisions/ Localities/ Urban Areas

Country name	Country code	Country	State	County	City
Iran	IRN	Country	Provinces / Ostanha	Counties/ Shahrestan	Localities/ Dehestân
Iraq	IRQ	Country	Governorates/ Muhafazat	Districts /Qadaa/ Kaza	Urban Areas/ Localities
Iceland	ISL	Country	Regions/ L andsvaedi	Municipalities/Sveitarfelog	Localities/Urban Areas
Israel	ISR	Country	Districts	Cities/ Local Councils	Localities/Urban Areas
Italy	ITA	Country	Regioni	Provincie	Communes/ Localities/Urban Areas
Jamaica	JAM	Country	Counties	Parishes	Constituencies/ Localities/ Urban Areas

Country name	Country code	Country	State	County	City
Jordan	JOR	Country	Governorates	Districts	Subdistricts/Cities
Japan	JPN	Country	Prefectures		Cities/Districts/Municipalities
Kazakhstan	KAZ	Country	Regions/Oblystar	Districts/Audandar	Towns/Kent/Localities/Urban Areas
Kenya	KEN	Country	Counties	Constituencies	Localities/Urban Areas/Suburbs
Kyrgyzstan	KGZ	Country	Regions/Oblasttar	Districts/Raions	Localities/Urban Areas
Cambodia	KHM	Country	Provinces/Khaet	Districts/Srök	Communes/Khum/Localities/Urban Areas

Country name	Country code	Country	State	County	City
Kiribati	KIR	Country	Districts	Island Councils	
Saint Kitts and Nevis	KNA	Country	Parishes	States	Localities/Urban Areas
South Korea	KOR	Country	Provinces /Do	Districts/Si/Gun	Localities/Urban Areas
Kuwait	KWT	Country	Governorates/Muhafazah	Areas/Min taqah	Cities/Communities
Laos	LAO	Country	Provinces / Khoueng	Districts/Muang	Localities/Urban Areas
Lebanon	LBN	Country	Governorates/Muhafazat	Districts/Qadaa	Municipalities/Localities/Urban Areas
Liberia	LBR	Country	Counties	Districts	Clans/Localities/Urban Areas

Country name	Country code	Country	State	County	City
Libya	LBY	Country	Districts/ Shabiya		Cities/Localities/ Urban Areas
Saint Lucia	LCA	Country	Districts/ Quarters		Localities/ Urban Areas
Liechtenstein	LIE	Country	Districts/ Bezirk	Municipalities/ Gemeinden	Localities/ Urban Areas
Sri Lanka	LKA	Country	Provinces	Districts	Divisional Secretariats/ Localities/ Urban Areas
Lesotho	LSO	Country	Districts	Constituencies	Community Councils/ Localities
Lithuania	LTU	Country	Apskritis	Savivaldybė	Seniūnija

Country name	Country code	Country	State	County	City
Luxembourg	LUX	Country	Cantons/ K antounen/ Kantone	Communes/ Gemengen/ Gemeinden	Localities/ Ortschaft/ Uertschaft/ Cities
Latvia	LVA	Country	Regions	Municipalities/ Novadi	Pilsētas/ Pagasti/ Localities/ Urban Areas
Macao	MAC	Country	Parishes	Districts	
Saint Martin	MAF	Country			Localities/ Urban Areas
Morocco	MAR	Country	Regions	Provinces/ Prefectures	Communes/ Localities/ Urban Areas
Monaco	MCO	Country	Communes	Wards/ Quartiers	

Country name	Country code	Country	State	County	City
Moldova	MDA	Country	Raion	Comuna	Localities/ Urban Areas
Madagascar	MDG	Country	Regions/ Faritra	Districts	Communes/ Localities/ Urban Areas
Maldives	MDV	Country	Atolls/Cities	Islands	
Mexico	MEX	Country	Estados	Municipios/ Delegaciones	Colonias/ Localities/ Urban Areas
Marshall Islands	MHL	Country	Municipalities		
Macedonia	MKD	Country	Statistical Regions	Opština	Localities/ Urban Areas
Mali	MLI	Country	Regions	Communes	Localities/ Urban Areas

Country name	Country code	Country	State	County	City
Malta	MLT	Country	Districts	Local Councils/ Kunsilli Lokali	Localities/ Urban Areas
Myanmar	MMR	Country	States/ Regions/ Union Territories	Districts	Townships/ Localities/ Urban Areas
Montenegro	MNE	Country	Opštine/ Municipalities		Localities/ Urban Areas
Mongolia	MNG	Country	Regions	Provinces/ Aimags	Districts/ Sums/ Localities/ Urban Areas
Northern Mariana Islands	MNP	Country	Municipalities		
Mozambique	MOZ	Country	Provinces	Districts/ Distritos	Localities/ Urban Areas

Country name	Country code	Country	State	County	City
Mauritania	MRT	Country	Regions	Départements	Localities/Urban Areas
Montserrat	MSR	Country	Parishes	Regions	Localities/Urban Areas
Martinique	MTQ	Country	Arrondissements	Communes	Localities/Urban Areas
Mauritius	MUS	Country	Islands	Districts	Wards/ Localities/ Urban Areas
Malawi	MWI	Country	Regions	Districts	Localities/Urban Areas
Malaysia	MYS	Country	States/ Negeri	Districts/ Daïra/Daerah	Subdistricts/ Mukim/ Localities/ Urban Area/ Bahagian Kecil

Country name	Country code	Country	State	County	City
Mayotte	MYT	Country	Communes		Villages
Namibia	NAM	Country	Provinces	Constituencies	Suburbs/ Localities
New Caledonia	NCL	Country	Provinces	Communes	
Niger	NER	Country	Regions	Departments	Localities/ Urban Areas
Nigeria	NGA	Country	States	Local Government Areas	Towns/ Cities
Nicaragua	NIC	Country	Departments/ Departamentos	Municipalities/ Municipios	Localities/ Urban Areas
Niue	NIU	Country	Villages		Towns
Netherlands	NLD	Country	Counties/ Fylker	Districts/ Okonomisk	Municipalities, Kommuner, Localities, or Urban Areas

Country name	Country code	Country	State	County	City
Norway	NOR	Country	Counties/ Fylker	Districts/ Okonomisk	Municipal ities, Kommuner, Localities, or Urban Areas
Nepal	NPL	Country	Provinces / Pradeshaharu	Districts/ Jilla	Municipal ities/Loc alities/ Urban Areas
Nauru	NRU	Country	Districts		
New Zealand	NZL	Country	Regions	Territorial Authoriti es	Statistic al Areas/ Loc alities/ Urban Areas
Oman	OMN	Country	Governorates/ Muhafazah	Provinces /Wilayat	Cities/ Urban Areas/ Com munities
Pakistan	PAK	Country	Provinces	Districts	Localities/ Tehsils

Country name	Country code	Country	State	County	City
Panama	PAN	Country	Provinces /Provinci as	Districts/ Distrito	Corregimi entos/ Loc alities/ Urban Areas
Pitcairn Islands	PCN	Country	Islands		
Peru	PER	Country	Regions	Districts	Distritos /Localiti es/ Urban Areas
Philippines	PHL	Country	Regions/ Rehiyon	Provinces / Lalawiga n	Municipal ities/ Mun isipiyos/ Cities/Lu ngsod
Palau	PLW	Country	States		
Papua New Guinea	PNG	Country	Regions	Provinces	Districts /Localiti es/ Urban Areas

Country name	Country code	Country	State	County	City
Poland	POL	Country	Provinces / Voivodes hips	Counties/ Powiats	Communes/ Gminas/ Towns/ Dziel nicas
North Korea	PRK	Country	Provinces		Localitie s/Urban Areas
Portugal	PRT	Country	Districts/ Distritos	Municipal ities/Con celhos	Civil Parish/Fr eguesias/ Localitie s/Urban Areas
Paraguay	PRY	Country	Departmer ts	Distritos	Localitie s/Urban Areas
Palestine	PSE	Country	Territori es	Governora tes/ Muhaf azat	Localitie s/Urban Areas
French Polynesia	PYF	Country	Subdivisi ons/lles	Communes	

Country name	Country code	Country	State	County	City
Qatar	QAT	Country	Municipal ities/Bal adiyat	Zones	Localitie s/Urban Areas
Réunion	REU	Country	Arrondiss ements	Communes	Localitie s/Urban Areas
Romania	ROU	Country	Regions/ Judete	Communes	Towns/ Oraş
Russia	RUS	Country	Federal District/ Federal'n yy Okrug	Oblast'	Rayon/ Raion/ Urban Area/ Gorod
Rwanda	RWA	Country	Provinces	Districts	Sectors/ S ecteurs/ L ocalities /Urban Areas
Saudi Arabia	SAU	Country	Regions/ Manatiq	Governora tes/ Muhaf azat	Municipal ities/ Amanah

Country name	Country code	Country	State	County	City
Sudan	SDN	Country	States/ Wilaya'at		Localities/ Urban Areas
Senegal	SEN	Country	Regions	Departments	Arrondissements/ Localities/ Urban Areas
Singapore	SGP	Country	Districts	Constituencies	Wards
Saint Helena	SHN	Country	Islands	Districts	Localities/ Urban Areas
Solomon Islands	SLB	Country	Provinces	Constituencies	Wards
Sierra Leone	SLE	Country	Provinces	Districts	Chiefdoms/ Localities/ Urban Areas
El Salvador	SLV	Country	Departments/ Departamentos	Municipalities/ Municipios	Localities/ Urban Areas

Country name	Country code	Country	State	County	City
San Marino	SMR	Country	Municipalities/Castelli	Localities/Urban Areas	
Somalia	SOM	Country	Regions/Globalada		Localities/Urban Areas
Saint Pierre and Miquelon	SPM	Country	Communes		
Serbia	SRB	Country	Autonomous Pokrajina/Regions	Okrug/Districts	Opština/Municipalities/Localities/Urban Areas
South Sudan	SSD	Country	States/Wilayat	Counties	Localities/Urban Areas
São Tomé and Príncipe	STP	Country	Provinces	Districts	Localities/Urban Areas
Suriname	SUR	Country	Districts/Distrikt	Resorts	Localities/Urban Areas

Country name	Country code	Country	State	County	City
Slovakia	SVK	Country	Regions/ Kraje	Districts/ Okresy	Municipal ities/ Obec/ Mestská cast
Slovenia	SVN	Country	Regions/ Regi	Upravne Enote	Municipal ities/ Obcine/ Local ities/Urb an Areas
Sweden	SWE	Country	Counties	Municipal ities	Localitie s/Urban Areas
Eswatini	SWZ	Country	Regions	Tinkhundl a	Towns/ Suburbs/ Loca lities
Sint Maarten	SXM	Country	Settlemen ts		
Seychelles	SYC	Country	Districts		Localitie s/Urban Areas

Country name	Country code	Country	State	County	City
Syria	SYR	Country	Governorates	Districts/Muhafazah	Cities/Localities/Urban Areas
Turks and Caicos Islands	TCA	Country	Districts	Localities	
Chad	TCD	Country	Regions	Départements	Arrondissements/Localities/Urban Areas
Togo	TGO	Country	Regions/Provinces	Prefectures	Localities/Urban Areas
Thailand	THA	Country	Provinces / Changwat	Districts/ Amphoe	Subdistricts/ Tambon/Localities/Urban Areas
Tajikistan	TJK	Country	Provinces /Regions	Districts/ Raion/Raion	Localities/Urban Areas

Country name	Country code	Country	State	County	City
Tokelau	TKL	Country	Atolls		
Turkmenistan	TKM	Country	Provinces /Welayat	Districts/ Etraplar	Towns
East Timor (Timor-Le ste)	TLS	Country	Municipal ities	Administr ative Post	Localitie s/Urban Areas
Tonga	TON	Country	Subdivisi ons		
Trinidad and Tobago	TTO	Country	Municipal ities		Localitie s/Urban Areas
Tunisia	TUN	Country	Governate s/ Wilayahs	Delegatio ns/ Mutama diyats	Municipal ities/Sh aykhats/ Lo calities/ Urban Areas

Country name	Country code	Country	State	County	City
Turkey	TUR	Country	Provinces /Il	Districts/ Ilce	Urban Areas/ Belde/ Subdistri cts/ Bucak/ Neighbor hoods/ Mahalle
Tuvalu	TUV	Country	Islands		
Taiwan	TWN	Country	Provinces	Counties	Townships /Local Neighborh oods
Tanzania	TZA	Country	Provinces /Mkoa	Districts/ Wilaya	Localitie s/Urban Areas
Uganda	UGA	Country	Regions	Districts	Counties/ Localitie s/Urban Areas

Country name	Country code	Country	State	County	City
Ukraine	UKR	Country	Oblast/ Mista/ Avton omna Respublik a	Raions	Settlemen t Councils/ Rural Councils/ Localitie s/Urban Areas
United States Minor Outlying Islands	UMI	Country	Islands/ Atolls		
Uruguay	URY	Country	Departmer ts/ Depart amentos	Municipio s/ Municip alities/S ecciones	Segmentos /Localiti es/ Urban Areas
United States of America	USA	Country	States/ Te rritories	Counties	MCD/ CCD/ Post Localitie s/ Municip alities
Uzbekistan	UZB	Country	Regions/ V iloyatlar	Districts/ Tumanlar	Localitie s/Urban Areas

Country name	Country code	Country	State	County	City
Vatican City	VAT	Country			Localities/Urban Areas
Saint Vincent and the Grenadines	VCT	Country	Parishes	Divisions	Localities/Urban Areas
Venezuela	VEN	Country	States/ Estados	Municipalities/ Municipios	Localities/Urban Areas/ Parish/ Parroquias
British Virgin Islands	VGB	Country	Districts		
Vietnam	VNM	Country	Provinces/ Cities	Districts	Wards/ Localities/ Urban Areas
Vanuatu	VUT	Country	Provinces		
Wallis and Futuna Islands	WLF	Country	Districts/ Rayaumes		
Samoa	WSM	Country	Districts/ Itūmālō	Towns	Localities/Urban Areas

Country name	Country code	Country	State	County	City
Kosovo	XKS	Country	Districts	Municipalities	Localities/Urban Areas
Yemen	YEM	Country	Governorates/ Muhafazat	Districts/ Muderiah	Localities/Urban Areas
South Africa	ZAF	Country	Provinces	Districts	Municipalities/Wards
Zambia	ZMB	Country	Provinces	Districts	Suburbs/ Localities
Zimbabwe	ZWE	Country	Provinces	Districts/ Muderiah	Localities/Urban Areas

The following is a list of the supported postal code formats by country, including the number of digits and an example postal code.

Note

PO BOX zipcodes are not supported postal code formats. Union territory zip codes that are used in India are also not supported.

Supported postal codes

Country	Postal format	Example
Afghanistan	4 digit	1001
Albania	4 digit	1001
Algeria	5 digit	01000
American Samoa	5 digit	96799
Andorra	5 digit	AD100
Anguilla	6 digit	AI-2640
Argentina	5 digit	A4126
Armenia	2 digit	00
Australia	4 digit	0800
Austria	4 digit	1010
Azerbaijan	2 digit	01
Brunei Darussalam	6 digit	BA1111
Bahrain	4 digit	0101
Bangladesh	2 digit	10

Country	Postal format	Example
Belarus	6 digit	202115
Belgium	4 digit	1000
Bermuda	4 digit	CR 01
Bhutan	2 digit	11
Bosnia and Herzegovina	5 digit	70101
Brazil	5 digit	01001
British Indian Ocean Territory	Alphanumeric - 5 digit	BBND 1
British Virgin Islands	4 digit	1110
Bulgaria	4 digit	1000
Cabo Verde	4 digit	1101
Cambodia	2 digit	01
Canada	3 digit	A0A
Cayman Islands	Alphanumeric - 7 digit	KY1-1000
Chile	3 digit	100
China	4 digit	0100

Country	Postal format	Example
Colombia	4 digit	0500
Costa Rica	5 digit	10101
Croatia	5 digit	10000
Cuba	1 digit	1
Cyprus	4 digit	1010
Czechia	5 digit	100 00
Democratic Republic of the Congo	4 digit	1001
Denmark	4 digit	1050
Dominican Republic	5 digit	10101
Ecuador	6 digit	010101
Egypt	2 digit	11
El Salvador	4 digit	1101
Estonia	5 digit	10001
Falkland Islands	Alphanumeric - 5 digit	FIQQ 1

Country	Postal format	Example
Faroe Islands	3 digit	100
Finland	5 digit	00100
France	5 digit	01000
French Guiana	5 digit	97300
French Polynesia	5 digit	98701
Georgia	2 digit	01
Germany	5 digit	01067
Ghana	2 digit	A2
Gibraltar	Alphanumeric - 5 digit	GX11 1
Greece	5 digit	104 31
Greenland	4 digit	3900
Guadeloupe	5 digit	97100
Guam	5 digit	96910
Guatemala	5 digit	01001

Country	Postal format	Example
Guernsey	Alphanumeric - 4 digit, 5 digit	GY1 1, GY10 1
Guinea-Bissau	4 digit	1000
Haiti	4 digit	1110
Holy See	5 digit	00120
Honduras	2 digit	11
Hungary	4 digit	1007
Iceland	3 digit	101
India	6 digit	110001
Indonesia	5 digit	10110
Iran	2 digit	11
Iraq	2 digit	10
Ireland	3 digit	A41
Isle of Man	Alphanumeric - 4 digit	IM1 1
Israel	5 digit	10292

Country	Postal format	Example
Italy	5 digit	00010
Japan	7 digit	001-0010
Jersey	Alphanumeric - 4 digit	JE2 3
Jordan	5 digit	11100
Kazakhstan	4 digit	0100
Kenya	1 digit	0
Kiribati	6 digit	KI0101
Kosovo	5 digit	10000
Kuwait	2 digit	00
Kyrgyzstan	4 digit	7200
Laos	2 digit	01
Latvia	4 digit	1001
Lesotho	1 digit	1
Liberia	2 digit	10

Country	Postal format	Example
Liechtenstein	4 digit	9485
Lithuania	5 digit	00100
Luxembourg	4 digit	1110
Macedonia	4 digit	1000
Madagascar	3 digit	101
Malawi	3 digit	101
Malaysia	5 digit	01000
Maldives	2 digit	00
Malta	3 digit	ATD
Marshall Islands	3 digit	969
Martinique	5 digit	97200
Mauritius	3 digit	111
Mayotte	5 digit	97600
Mexico	5 digit	01000

Country	Postal format	Example
Micronesia	5 digit	96941
Moldova	4 digit	2001
Monaco	5 digit	98000
Mongolia	4 digit	1200
Montenegro	5 digit	81000
Montserrat	4 digit	1120
Morocco	5 digit	10000
Mozambique	4 digit	1100
Myanmar	2 digit	01
Namibia	3 digit	100
Nepal	3 digit	101
Netherlands	4 digit	1011
New Caledonia	5 digit	98800
New Zealand	4 digit	0110

Country	Postal format	Example
Nicaragua	3 digit	110
Niger	4 digit	1000
Nigeria	4 digit	1002
Niue	4 digit	9974
Norfolk Island	4 digit	2899
Northern Mariana Islands	5 digit	96950
Norway	4 digit	0010
Oman	1 digit	1
Pakistan	2 digit	10
Palau	5 digit	96939
Palestine	4 digit	P104
Papua New Guinea	3 digit	111
Paraguay	6 digit	001001
Peru	5 digit	01000

Country	Postal format	Example
Philippines	4 digit	1000
Pitcairn	Alphanumeric - 5 digit	PCRN 1
Poland	5 digit	00-002
Portugal	4 digit	1000
Puerto Rico	5 digit	00601
Romania	6 digit	010011
Russia	6 digit	101000
Réunion	5 digit	97400
Saint Barthélemy	5 digit	97133
Saint Helena, Ascension and Tristan da Cunha	Alphanumeric - 5 digit	ASCN 1
Saint Lucia	7 digit	LC01 101
Saint Martin	5 digit	97150
Saint Pierre and Miquelon	5 digit	97500
Saint Vincent and the Grenadines	4 digit	VC01

Country	Postal format	Example
Samoa	2 digit	11
San Marino	5 digit	47890
Saudi Arabia	2 digit	12
Senegal	5 digit	10000
Serbia	5 digit	11000
Singapore	6 digit	018906
Slovakia	5 digit	010 01
Slovenia	4 digit	1000
South Africa	4 digit	0001
South Georgia and the South Sandwich Islands	Alphanumeric - 5 digit	SIQQ 1
South Korea	5 digit	01000
Spain	5 digit	01001
Sri Lanka	2 digit	00
Sudan	2 digit	11

Country	Postal format	Example
Svalbard and Jan Mayen	4 digit	8099
Swaziland	1 digit	H
Sweden	5 digit	111 15
Switzerland	4 digit	1000
Taiwan	3 digit	100
Tajikistan	4 digit	7340
Tanzania, United Republic of	3 digit	111
Thailand	5 digit	10100
Timor-Leste	4 digit	TL10
Trinidad and Tobago	2 digit	10
Tunisia	4 digit	1000
Turkey	5 digit	01010
Turkmenistan	3 digit	744
Turks and Caicos Islands	Alphanumeric - 5 digit	TKCA 1

Country	Postal format	Example
U.S. Virgin Islands	5 digit	00802
Ukraine	3 digit, 5 digit	070, 01001
United Kingdom	Alphanumeric - 2 to 5 digits	B1, AL1, AB10, AB10 1
United States	5 digit	00001
Uruguay	5 digit	11000
Uzbekistan	4 digit	1000
Venezuela	4 digit	0000
Vietnam	5 digit	01106
Wallis and Futuna	5 digit	98600
Zambia	5 digit	10100

Using unsupported or custom dates

Amazon Quick Sight natively supports a limited number of date formats. However, you can't always control the format of the data provided to you. When your data contains a date in an unsupported format, you can tell Amazon Quick Sight how to interpret it.

You can do this by editing the dataset, and changing the format of the column from text or numeric to date. A screen appears after you make this change, so you can enter the format. For example, if you are using a relational data source, you can specify MM-dd-yyyy for a text field

containing '09-19-2017', so it is interpreted as 2017-09-19T00:00:00.000Z. If you are using a nonrelational data source, you can do the same thing starting with a numeric field or a text field.

Amazon Quick Sight only supports text to date for relational (SQL) sources.

For more information on supported date formats, see [Supported date formats](#).

Use this procedure to help Amazon Quick Sight understand dates in different formats.

1. For a dataset containing unsupported date formats, edit the data as follows. For the column containing your datetime data, change the data type from text to date. Do this by choosing the colorful data type icon beneath the column name in the data preview.

 Note

Integer dates that aren't Unix epoch datetimes don't work as is. For example, these formats are not supported as integers: MMddyy, MMddyyyy, ddMMyy, ddMMyyyy, and yyMMdd. The workaround is to first change them to text format. Make sure all your rows contain six digits (not five). Then change the text data type to datetime. For more information on Unix epoch datetimes, see [epochDate](#).

When you change the data type to date, the **Edit date format** screen appears.

2. Enter your date format, indicating which parts are month, date, year, or time. Formats are case-sensitive.
3. Choose **Validate** to make sure Amazon Quick Sight can now interpret your datetime data with the format you specified. Rows that don't validate are skipped and omitted from the dataset.
4. When you are satisfied with the results, choose **Update**. Otherwise, choose **Close**.

Adding a unique key to an Amazon Quick Sight dataset

Quick authors can configure a unique key column to a Quick Sight dataset during data preparation. This unique key acts as a global sorting key for the dataset and optimizes query generation for table visuals. When a user creates a table visual in Quick Sight and adds the unique key column to the value field well, data is sorted from left to right up to the unique key column. All columns to the right of the unique key column are ignored in the sort order. Tables that don't contain a unique key are sorted based on the order that the columns appear in the dataset.

The following limitations apply to unique keys:

- Unique keys are only supported for unaggregated tables.
- If a dataset column is used for column level security (CLS), the column can't also be used as the unique key.

Use the following procedure to designate a unique key for a dataset in Amazon Quick Sight.

To set up a unique key

1. Open the [Quick console](#).
2. Choose **Data**.
3. Perform one of the following actions:
 - a. Navigate to the dataset that you want to add a unique key to, choose the ellipsis (three dots) next to the dataset, and then choose **Edit**.
 - b. Choose **New** then **Dataset**. Choose the dataset that you want to add and then choose **Edit data source**. For more information about creating new datasets in Amazon Quick Sight, see [Creating datasets](#).
4. The data preparation page for the dataset opens. Navigate to the **Fields** pane and locate the field that you want to set as the unique key.
5. Choose the ellipsis (three dots) next to the field name, and then choose **Set as unique key**.

After you create a unique key, a key icon appears next to the field to show that the field is now the unique key for the dataset. When you save and publish the dataset, the unique key configuration is applied to the dataset and to all dashboards and analyses that are created with that dataset. To remove a unique key from a dataset, navigate to the data preparation page for the dataset, choose the ellipsis next to the unique key field, and then choose **Remove as unique key**. After you remove a unique key from a dataset, you can designate a different field as the unique key.

Integrating Amazon SageMaker AI models with Amazon Quick Sight

Note

You don't need any technical experience in machine learning (ML) to author analyses and dashboards that use the ML-powered features in Amazon Quick Sight.

You can augment your Amazon Quick Enterprise edition data with Amazon SageMaker AI machine learning models. You can run inferences on data stored in SPICE imported from any data source supported by Quick. For a full list of supported data sources, see [Supported data sources](#).

Using Quick with SageMaker AI models can save the time that you might otherwise spend managing data movement and writing code. The results are useful both for evaluating the model and—when you're satisfied with the results—for sharing with decision-makers. You can begin immediately after the model is built. Doing this surfaces your data scientists' prebuilt models, and enables you to apply the data science to your datasets. Then you can share these insights in your predictive dashboards. With the Quick serverless approach, the process scales seamlessly, so you don't need to worry about inference or query capacity.

Amazon Quick supports SageMaker AI models that use regression and classification algorithms. You can apply this feature to get predictions for just about any business use case. Some examples include predicting the likelihood of customer churn, employee attrition, scoring sales leads, and assessing credit risks. To use Quick to provide predictions, the SageMaker AI model data for both input and output must be in tabular format. In multiclass or multilabel classification use cases, each output column has to contain a single value. Quick doesn't support multiple values inside a single column.

Topics

- [How SageMaker AI integration works](#)
- [Costs incurred \(no additional costs with integration itself\)](#)
- [Usage guidelines](#)
- [Defining the schema file](#)
- [Adding a SageMaker AI model to your Quick Sight dataset](#)
- [Build predictive models with SageMaker AI Canvas](#)

How SageMaker AI integration works

In general, the process works like this:

1. An Amazon Quick administrator adds permissions for Quick to access SageMaker AI. To do this, open **Security & Permissions** settings from the **Manage Quick** page. Go to **Quick access to AWS services**, and add SageMaker AI.

When you add these permissions, Quick is added to an AWS Identity and Access Management (IAM) role that provides access to list all the SageMaker AI models in your AWS account. It also provides permissions to run SageMaker AI jobs that have names that are prefixed with `quicksight-auto-generated-`.

2. We recommend that you connect to an SageMaker AI model that has an inference pipeline, because it automatically performs data preprocessing. For more information, see [Deploy an Inference Pipeline](#) in the *SageMaker AI Developer Guide*.
3. After you identify the data and the pretrained model that you want to use together, the owner of the model creates and provides a schema file. This JSON file is a contract with SageMaker AI. It provides metadata about the fields, data types, column order, output, and settings that the model expects. The optional settings component provides the instance size and count of the compute instances to use for the job.

If you're the data scientist who built the model, create this schema file using the format documented following. If you're a consumer of the model, get the schema file from the owner of the model.

4. In Quick, you begin by creating a new dataset with the data that you want to make predictions on. If you're uploading a file, you can add the SageMaker AI model on the upload settings screen. Otherwise, add the model on the data preparation page.

Before you proceed, verify the mappings between the dataset and the model.

5. After the data is imported into the dataset, the output fields contain the data returned from SageMaker AI. You use these fields just as you use other fields, within the guidelines described in [Usage guidelines](#).

When you run SageMaker AI integration, Quick passes a request to SageMaker AI to run batch transform jobs with inference pipelines. Quick starts provisions and deployment of the instances needed in your AWS account. When processing is complete, these instances are shut down and terminated. The compute capacity incurs costs only when it's processing models.

To make it easier for you to identify them, Quick names all its SageMaker AI jobs with the prefix `quicksight-auto-generated-`.

6. The output of the inference is stored in SPICE and appended to the dataset. As soon as the inference is complete, you can use the dataset to create visualizations and dashboards using the prediction data.

7. The data refresh starts every time you save the dataset. You can start the data refresh process manually by refreshing the SPICE dataset, or you can schedule it to run at a regular interval. During each data refresh, the system automatically calls SageMaker AI batch transform to update the output fields with new data.

You can use the Amazon Quick Sight SPICE ingestion API operations to control the data refresh process. For more information about using these API operations, see the [Amazon Quick Sight API Reference](#).

Costs incurred (no additional costs with integration itself)

Using this feature doesn't require an additional fee in itself. Your costs include the following:

- The cost of model deployment through SageMaker AI, which is incurred only when the model is running. Saving a dataset—after either creating or editing it—or refreshing its data starts the data ingestion process. This process includes calling SageMaker AI if the dataset has inferred fields. Costs are incurred in the same AWS account where your Quick subscription is.
- Your Quick subscription costs are as follows:
 - The cost of storing your data in the in-memory calculation engine in Quick (SPICE). If you are adding new data to SPICE, you might need to purchase enough SPICE capacity to accommodate it.
 - Quick subscriptions for the authors or admins who build the datasets.
 - Pay-per-session charges for viewers (readers) to access interactive dashboards.

Usage guidelines

In Amazon Quick, the following usage guidelines apply to this Enterprise edition feature:

- The processing of the model occurs in SPICE. Therefore, it can only apply to datasets that are stored in SPICE. The process currently supports up to 500 million rows per dataset.
- Only Quick admins or authors can augment datasets with ML models. Readers can only view the results when they are part of a dashboard.
- Each dataset can work with one and only one ML model.
- Output fields can't be used to calculate new fields.
- Datasets can't be filtered by fields that are integrated with the model. In other words, if your dataset field is currently mapped to the ML model, you can't filter on that field.

In SageMaker AI, the following usage guidelines apply to a pretrained model that you use with Amazon Quick Sight:

- When you create the model, associate it with the Amazon Resource Name (ARN) for the appropriate IAM role. The IAM role for the SageMaker AI model needs to have access to the Amazon S3 bucket that Amazon Quick Sight uses.
- Make sure that your model supports .csv files for both input and output. Make sure that your data is in a tabular format.
- Provide a schema file that contains metadata about the model, including the list of input and output fields. Currently, you must create this schema file manually.
- Consider the amount of time that it takes to complete your inference, which depends on a number of factors. These include the complexity of the model, the amount of data, and the compute capacity defined. Completing the inference can take several minutes to several hours. Amazon Quick Sight caps all data ingestion and inferencing jobs to a maximum of 10 hours. To reduce the time it takes to perform an inference, consider increasing the instance size or the number of instances.
- Currently, you can use only batch transforms for integration with SageMaker AI, not real-time data. You can't use an SageMaker AI endpoint.

Defining the schema file

Before you use an SageMaker AI model with Quick Sight data, create the JSON schema file that contains the metadata that Amazon Quick Sight needs to process the model. The Amazon Quick author or admin uploads the schema file when configuring the dataset.

The schema fields are defined as follows. All fields are required unless specified in the following description. Attributes are case-sensitive.

inputContentType

The content type that this SageMaker AI model expects for the input data. The only supported value for this is "text/csv". Quick Sight doesn't include any of the header names that you add to the input file.

outputContentType

The content type of the output that is produced by the SageMaker AI model that you want to use. The only supported value for this is "text/csv".

input

A list of features that the model expects in the input data. Quick Sight produces the input data in exactly the same order. This list contains the following attributes:

- *name* – The name of the column. If possible, make this the same as the name of the corresponding column in the QuickSight dataset. This attribute is limited to 100 characters.
- *type* – The data type of this column. This attribute takes the values "INTEGER", "STRING", and "DECIMAL".
- *nullable* – (Optional) The nullability of the field. The default value is `true`. If you set `nullable` to `false`, Quick Sight drops rows that don't contain this value before calling SageMaker AI. Doing this helps avoid causing SageMaker AI to fail on missing required data.

output

A list of output columns that the SageMaker AI model produces. Quick Sight expects these fields in exactly the same order. This list contains the following attributes:

- *name* – This name becomes the default name for the corresponding new column that's created in Quick Sight. You can override the name specified here in Quick Sight. This attribute is limited to 100 characters.
- *type* – The data type of this column. This attribute takes the values "INTEGER", "STRING", and "DECIMAL".

instanceTypes

A list of the ML instance types that SageMaker AI can provision to run the transform job. The list is provided to the Amazon Quick user to choose from. This list is limited to the types supported by SageMaker AI. For more information on supported types, see [TransformResources](#) in the *SageMaker AI Developer Guide*.

defaultInstanceType

(Optional) The instance type that is presented as the default option in the SageMaker AI wizard in Quick Sight. Include this instance type in `instanceTypes`.

instanceCount

(Optional) The instance count defines how many of the selected instances for SageMaker AI to provision to run the transform job. This value must be a positive integer.

description

This field provides a place for the person who owns the SageMaker AI model to communicate with the person who is using this model in Quick Sight. Use this field to provide hints about successfully using this model. For example, this field can contain information about selecting an effective instance type to choose from the list in `instanceTypes`, based on the size of dataset. This field is limited to 1,000 characters.

version

The version of the schema, for example "1.0".

The following example shows the structure of the JSON in the schema file.

```
{
  "inputContentType": "CSV",
  "outputContentType": "CSV",
  "input": [
    {
      "name": "buying",
      "type": "STRING"
    },
    {
      "name": "maint",
      "type": "STRING"
    },
    {
      "name": "doors",
      "type": "INTEGER"
    },
    {
      "name": "persons",
      "type": "INTEGER"
    },
    {
      "name": "lug_boot",
      "type": "STRING"
    },
    {
      "name": "safety",
      "type": "STRING"
    }
  ],
}
```



```
    "output": [
      {
        "name": "Acceptability",
        "type": "STRING"
      }
    ],
    "description": "Use ml.m4.xlarge instance for small datasets, and ml.m4.4xlarge
for datasets over 10 GB",
    "version": "1.0",
    "instanceCount": 1,
    "instanceTypes": [
      "ml.m4.xlarge",
      "ml.m4.4xlarge"
    ],
    "defaultInstanceType": "ml.m4.xlarge"
  }
}
```

The structure of the schema file is related to the kind of model that is used in examples provided by SageMaker AI.

Adding a SageMaker AI model to your Quick Sight dataset

Using the following procedure, you can add a pretrained SageMaker AI model to your dataset, so that you can use predictive data in analyses and dashboards.

Before you begin, have the following items available:

- The data that you want to use to build the dataset.
- The name of the SageMaker AI model that you want to use to augment the dataset.
- The schema of the model. This schema includes field name mappings and data types. It's helpful if it also contains recommended settings for instance type and number of instances to use.

To augment your Amazon Quick Sight dataset with SageMaker AI

1. Create a new dataset from the start page by choosing **Datasets**, and then choose **New dataset**.

You can also edit an existing dataset.

2. Choose **Augment with SageMaker** on the data preparation screen.
3. For **Select your model**, choose the following settings:

- **Model** – Choose the SageMaker AI model to use to infer fields.

- **Name** – Provide a descriptive name for the model.
- **Schema** – Upload the JSON schema file provided for the model.
- **Advanced settings** – QuickSight recommends the selected defaults based on your dataset. You can use specific runtime settings to balance the speed and cost of your job. To do this, enter the SageMaker AI ML instance types for **Instance type** and number of instances for **Count**.

Choose **Next** to continue.

4. For **Review inputs**, review the fields that are mapped to your dataset. Quick Sight attempts to automatically map the fields in your schema to the fields in your dataset. You can make changes here if the mapping needs adjustment.

Choose **Next** to continue.

5. For **Review outputs**, view the fields that are added to your dataset.

Choose **Save and prepare data** to confirm your choices.

6. To refresh the data, choose the dataset to view details. Then either choose **Refresh Now** to manually refresh the data, or choose **Schedule refresh** to set up a regular refresh interval. During each data refresh, the system automatically runs the SageMaker AI batch transform job to update the output fields with new data.

Build predictive models with SageMaker AI Canvas

Amazon Quick authors can export data into SageMaker AI Canvas to build ML models that can be sent back to Quick. Authors can use these ML models to augment their datasets with predictive analytics that can be used to build analyses and dashboards.

Prerequisites

- A Quick account that's integrated with IAM Identity Center. If your Quick account isn't integrated with IAM Identity Center, create a new Quick account and choose **Use IAM Identity Center enabled application** as the identity provider.
 - For more information on IAM Identity Center, see [Getting started](#).
 - To learn more about integrating your Quick with IAM Identity Center, see [Configure your Amazon Quick account with IAM Identity Center](#).

- To import assets from an existing Quick account to a new Quick account that's integrated with IAM Identity Center, see [Asset bundle operations](#).
- A new SageMaker AI domain that is integrated with IAM Identity Center. For more information about onboarding to SageMaker AI Domain with IAM Identity Center, see [Onboard to SageMaker AI Domain using IAM Identity Center](#).

Topics

- [Build a predictive model in SageMaker AI Canvas from Amazon Quick Sight](#)
- [Create a dataset with a SageMaker AI Canvas model](#)
- [Considerations](#)

Build a predictive model in SageMaker AI Canvas from Amazon Quick Sight

To build a predictive model in SageMaker AI Canvas

1. Log in to Amazon Quick and navigate to the tabular table or pivot table that you want to create a predictive model for.
2. Open the on-visual menu and choose **Build a predictive model**.
3. In the **Build a predictive model in SageMaker AI Canvas** pop up that appears, review the information presented and then choose **EXPORT DATA TO SAGEMAKER CANVAS**.
4. In the **Exports** pane that appears, choose **GO TO SAGEMAKER CANVAS** when the export is completed to go to the SageMaker AI Canvas console.
5. In SageMaker AI Canvas, create a predictive model with the data that you exported from Quick Sight. You can choose to follow a guided tour that helps you create the predictive model, or you can skip the tour and work at your own pace. For more information about creating a predictive model in SageMaker AI Canvas, see [Build a model](#).
6. Send the predictive model back to Quick Sight. For more information about sending a model from SageMaker AI Canvas to Amazon Quick Sight, see [Send your model to Amazon Quick Sight](#).

Create a dataset with a SageMaker AI Canvas model

After you create a predictive model in SageMaker AI Canvas and send it back to Quick Sight, use the new model to create a new dataset or apply it to an existing dataset.

To add a predictive field to a dataset

1. Open the Quick console, choose **Data** at left, and choose the **Datasets** tab.
2. Upload a new dataset or choose an existing dataset.
3. Choose **Edit**.
4. On the dataset's data prep page, choose **ADD**, and then choose **Add predictive field** to open the **Augment with SageMaker AI** modal.
5. For **Model**, choose the model that you sent to Quick Sight from SageMaker AI Canvas. The schema file automatically populates in the **Advanced settings** pane. Review the inputs, and then choose **Next**.
6. On the **Review outputs** pane, enter a field name and description for a column to be targeted by the model that you created in SageMaker AI Canvas.
7. When you are finished, choose **Prepare data**.
8. After you choose **Prepare data**, you are redirected to the dataset page. To publish the new dataset, choose, **Publish & Visualize**.

When you publish a new dataset that uses a model from SageMaker AI Canvas, the data is imported into SPICE and a batch inference job begins in SageMaker AI. It can take up to 10 minutes for these processes to complete.

Considerations

The following limitations apply to the creation of SageMaker AI Canvas models with Quick Sight data.

- The **Build a predictive model** option that is used to send data to SageMaker AI Canvas is only available on table and tabular pivot table visuals. The table or pivot table visual must have between 2 and 1,000 fields and at least 500 rows.
- Datasets that contain integer or geographic data types will experience schema mapping errors when you add a predictive field to the dataset. To resolve this issue, remove the integer or geographic data types from the dataset or convert them to a new data type.

Preparing dataset examples

You can prepare data in any dataset to make it more suitable for analysis, for example changing a field name or adding a calculated field. For database datasets, you can also determine the data used by specifying a SQL query or joining two or more tables.

Use the following topics to learn how to prepare datasets.

Topics

- [Preparing a dataset based on file data](#)
- [Preparing a dataset based on Salesforce data](#)
- [Preparing a dataset based on database data](#)

Preparing a dataset based on file data

Use the following procedure to prepare a dataset based on text or Microsoft Excel files from either your local network or Amazon S3.

To prepare a dataset based on text or Microsoft Excel files from a local network or S3

1. Open a file dataset for data preparation by choosing one of the following options:
 - Create a new local file dataset, and then choose **Edit/Preview data**. For more information about creating a new dataset from a local text file, see [Creating a dataset using a local text file](#). For more information about creating a new dataset from a Microsoft Excel file, see [Creating a dataset using a Microsoft Excel file](#).
 - Create a new Amazon S3 dataset, and then choose **Edit/Preview data**. For more information about creating a new Amazon S3 dataset using a new Amazon S3 data source, see [Creating a dataset using Amazon S3 files](#). For more information about creating a new Amazon S3 dataset using an existing Amazon S3 data source, see [Creating a dataset using an existing Amazon S3 data source](#).
 - Open an existing Amazon S3, text file, or Microsoft Excel dataset for editing, from either the analysis page or the **Your Datasets** page. For more information about opening an existing dataset for data preparation, see [Editing datasets](#).
2. (Optional) On the data preparation page, enter a new name into the dataset name box on the application bar.

This name defaults to the file name for local files. For example, it defaults to **Group 1** for Amazon S3 files.

3. Review the file upload settings and correct them if necessary. For more information about file upload settings, see [Choosing file upload settings](#).

 Important

If you want to change upload settings, make this change before you make any other changes to the dataset. New upload settings cause Amazon Quick Sight to reimport the file. This process overwrites all of your other changes.

4. Prepare the data by doing one or more of the following:
 - [Selecting fields](#)
 - [Editing field names and descriptions](#)
 - [Changing a field data type](#)
 - [Adding calculated fields](#)
 - [Filtering data in Amazon Quick Sight](#)
5. Check the [SPICE](#) indicator to see if you have enough capacity to import the dataset. File datasets automatically load into SPICE. The import happens when you choose either **Save & visualize** or **Save**.

If you don't have access to enough SPICE capacity, you can make the dataset smaller by using one of the following options:

- Apply a filter to limit the number of rows.
- Select fields to remove from the dataset.

 Note

The SPICE indicator doesn't update to how much space you save by removing fields or filtering the data. It continues to reflect the SPICE usage from the last import.

6. Choose **Save** to save your work, or **Cancel** to cancel it.

You might also see **Save & visualize**. This option appears based on the screen that you started from. If this option isn't there, you can create a new visualization by starting from the dataset screen.

Preparing a dataset based on a Microsoft Excel file

Use the following procedure to prepare a Microsoft Excel dataset.

To prepare a Microsoft Excel dataset

1. Open a text file dataset for preparation by choosing one of the following options:
 - Create a new Microsoft Excel dataset, and then choose **Edit/Preview data**. For more information about creating a new Excel dataset, see [Creating a dataset using a Microsoft Excel file](#).
 - Open an existing Excel dataset for editing. You can do this from the analysis page or the **Your Datasets** page. For more information about opening an existing dataset for data preparation, see [Editing datasets](#).
2. (Optional) On the data preparation page, enter a name into the dataset name box in the application bar. If you don't rename the dataset, its name defaults to the Excel file name.
3. Review the file upload settings and correct them if necessary. For more information about file upload settings, see [Choosing file upload settings](#).

Important

If it's necessary to change upload settings, make this change before you make any other changes to the dataset. Changing upload settings causes Amazon Quick Sight to reimport the file. This process overwrites any changes you have made so far.

4. (Optional) Change the worksheet selection.
5. (Optional) Change the range selection. To do this, open **Upload Settings** from the on-dataset menu beneath the login name at upper right.
6. Prepare the data by doing one or more of the following:
 - [Selecting fields](#)
 - [Editing field names and descriptions](#)

- [Changing a field data type](#)
 - [Adding calculated fields](#)
 - [Filtering data in Quick Sight](#)
7. Check the [SPICE](#) indicator to see if you have enough space to import the dataset. Amazon Quick Sight must import Excel datasets into SPICE. This import happens when you choose either **Save & visualize** or **Save**.

If you don't have enough SPICE capacity, you can choose to make the dataset smaller using one of the following methods:

- Apply a filter to limit the number of rows.
- Select fields to remove from the dataset.
- Define a smaller range of data to import.

 Note

The SPICE indicator doesn't update to reflect your changes until after you load them. It shows the SPICE usage from the last import.

8. Choose **Save** to save your work, or **Cancel** to cancel it.

You might also see **Save & visualize**. This option appears based on the screen that you started from. If this option isn't there, you can create a new visualization by starting from the dataset screen.

Preparing a dataset based on Salesforce data

Use the following procedure to prepare a Salesforce dataset.

To prepare a Salesforce dataset

1. Open a Salesforce dataset for preparation by choosing one of the following options:
 - Create a new Salesforce dataset and choose **Edit/Preview data**. For more information about creating a new Salesforce dataset using a new Salesforce data source, see [Creating a dataset from Salesforce](#). For more information about creating a new Salesforce dataset using an existing Salesforce data source, see [Create a dataset using an existing Salesforce data source](#).

- Open an existing Salesforce dataset for editing from either the analysis page or the **Your Datasets** page. For more information about opening an existing dataset for data preparation, see [Editing datasets](#).
2. (Optional) On the data preparation page, enter a name into the dataset name box in the application bar if you want to change the dataset name. This name defaults to the report or object name.
 3. (Optional) Change the data element selection to see either reports or objects.
 4. (Optional) Change the data selection to choose a different report or object.

If you have a long list in the **Data** pane, you can search to locate a specific item by entering a search term into the **Search tables** box. Any item whose name contains the search term is shown. Search is case-insensitive and wildcards are not supported. Choose the cancel icon (X) to the right of the search box to return to viewing all items.

5. Prepare the data by doing one or more of the following:
 - [Selecting fields](#)
 - [Editing field names and descriptions](#)
 - [Changing a field data type](#)
 - [Adding calculated fields](#)
 - [Filtering data in Quick Sight](#)
6. Check the [SPICE](#) indicator to see if you have enough space to import the dataset. Importing data into SPICE is required for Salesforce datasets. Importing occurs when you choose either **Save & visualize** or **Save**.

If you don't have enough SPICE capacity, you can remove fields from the dataset or apply a filter to decrease its size. For more information about adding and removing fields from a dataset, see [Selecting fields](#).

 Note

The SPICE indicator doesn't update to reflect the potential savings of removing fields or filtering the data. It continues to reflect the size of the dataset as retrieved from the data source.

7. Choose **Save** to save your work, or **Cancel** to cancel it.

You might also see **Save & visualize**. This option appears based on the screen you started from. If this option isn't there, you can create a new visualization by starting from the dataset screen.

Preparing a dataset based on database data

Use the following procedure to prepare a dataset based on a query to a database. The data for this dataset can be from an AWS database data source like Amazon Athena, Amazon RDS, or Amazon Redshift, or from an external database instance. You can choose whether to import a copy of your data into [SPICE](#), or to query the data directly.

To prepare a dataset based on a query to a database

1. Open a database dataset for preparation by choosing one of the following options:
 - Create a new database dataset and choose **Edit/Preview data**. For more information about creating a new dataset using a new database data source, see [Creating a dataset from a database](#). For more information about creating a new dataset using an existing database data source, see [Creating a dataset using an existing database data source](#).
 - Open an existing database dataset for editing from either the analysis page or the **Your Datasets** page. For more information about opening an existing dataset for data preparation, see [Editing datasets](#).
2. (Optional) On the data preparation page, enter a name into the dataset name box on the application bar.

This name defaults to the table name if you selected one before data preparation. Otherwise, it's **Untitled data source**.

3. Decide how your data is selected by choosing one of the following:
 - To use a single table to provide data, choose a table or change the table selection.

If you have a long table list in the **Tables** pane, you can search for a specific table by typing a search term for **Search tables**.

Any table whose name contains the search term is shown. Search is case-insensitive and wildcards are not supported. Choose the cancel icon (X) to the right of the search box to return to viewing all tables.

- To use two or more joined tables to provide data, choose two tables and join them using the join pane. You must import data into Quick Sight if you choose to use joined tables. For more information about joining data using the Amazon Quick Sight interface, see [Joining data](#).
- To use a custom SQL query to provide data in a new dataset, choose **Switch to Custom SQL** tool on the **Tables** pane. For more information, see [Using SQL to customize data](#).

To change the SQL query in an existing dataset, choose **Edit SQL** on the **Fields** pane to open the SQL pane and edit the query.

4. Prepare the data by doing one or more of the following:

- [Selecting fields](#)
- [Editing field names and descriptions](#)
- [Changing a field data type](#)
- [Adding calculated fields](#)
- [Filtering data in Quick Sight](#)

5. If you aren't joining tables, choose whether to query the database directly or to import the data into SPICE by selecting either the **Query** or **SPICE** radio button. We recommend using SPICE for enhanced performance.

If you want to use SPICE, check the SPICE indicator to see if you have enough space to import the dataset. Importing occurs when you choose either **Save & visualize** or **Save**.

If you don't have enough space, you can remove fields from the dataset or apply a filter to decrease its size.

 Note

The SPICE indicator doesn't update to reflect the potential savings of removing fields or filtering the data. It continues to reflect the size of the dataset as retrieved from the data source.

6. Choose **Save** to save your work, or **Cancel** to cancel it.

You might also see an option to **Save & visualize**. This option appears based on the screen you started from. If this option isn't there, you can create a new visualization by starting from the dataset screen.

Analyses and reports: Visualizing data in Amazon Quick Sight

Following, you can find descriptions of how to create and customize Amazon Quick Sight charts, arrange charts in a dashboard, and more.

Topics

- [Working with an analysis in Amazon Quick Sight](#)
- [Adding and managing sheets](#)
- [Working with interactive sheets in Amazon Quick Sight](#)
- [Working with pixel perfect reports in Amazon Quick Sight](#)
- [Working with items on sheets in Amazon Quick Sight analyses](#)
- [Using themes in Amazon Quick Sight](#)
- [Accessing Amazon Quick Sight using keyboard shortcuts](#)

Working with an analysis in Amazon Quick Sight

In Quick Sight, an analysis is the same thing as a dashboard, except that it can only be accessed by the authors you choose. You can keep it private, and make it as robust and detailed as you like. When and if you decide to publish it, the shared version of it is called a dashboard.

Use the following sections to learn how to interact with a Quick Sight analysis.

Topics

- [Starting an analysis in Quick Sight](#)
- [Generating an analysis with natural language prompts](#)
- [Adding a title and description to an analysis](#)
- [Sharing Quick Sight analyses](#)
- [Renaming an analysis](#)
- [Duplicating analyses](#)
- [Customize date and time values of an analysis](#)
- [The analysis menu](#)
- [Configure analysis settings](#)
- [Item limits for Amazon Quick Sight analyses in the Quick Sight APIs](#)
- [Saving changes to analyses](#)

- [Exporting data from Quick Sight analyses](#)
- [Deleting an analysis](#)

Starting an analysis in Quick Sight

In Quick Sight, you analyze and visualize your data in analyses. When you're finished, you can publish your analysis as a dashboard to share with others in your organization.

Use the following procedure to create a new analysis.

To create a new analysis

1. On the Amazon Quick start page, choose **Analyses**, and then choose **New analysis**.
2. Choose the dataset that you want to include in your new analysis, and then choose **USE IN ANALYSIS** in the top right.
3. In the **New sheet** pop-up that appears, choose the sheet type that you want. You can choose between an **Interactive sheet** and a **Pixel perfect report**. To create a pixel perfect report, you need the pixel perfect reports add-on for your account. For more information about pixel perfect reports, see [Working with pixel perfect reports in Amazon Quick Sight](#). For more information on sheets, see [Adding and managing sheets](#).
4. (Optional) If you choose **Interactive sheet**, follow these steps:
 - (Optional): Choose the type layout that you want for your interactive sheet. You can choose one of the following options:
 - Free-form
 - Tiled

The default option is **Free-form**.

For more information about interactive sheet layouts, see [Types of layout](#).

- Choose the canvas size that you want your sheet optimized for. You can choose one of the following options:
 - 1024px
 - 1280px
 - 1366px
 - 1600px

- 1920px

For more information on formatting interactive sheets, see [Working with interactive sheets in Amazon Quick Sight](#).

5. (Optional) If you choose **Pixel perfect report**, follow these steps:

- (Optional) Choose the paper size that you want for your paginated report. You can choose from the following options:
 - US letter (8.5 x 11 in.)
 - US legal (8.5 x 14 in.)
 - A0 (841 x 1189 mm)
 - A1 (594 x 841 mm)
 - A2 (420 x 594 mm)
 - A3 (297 x 420 mm)
 - A4 (210 x 297 mm)
 - A5 (148 x 210 mm)
 - Japan B4 (257 x 364 mm)
 - Japan B5 (182 x 257 mm)

The default paper size is US letter (8.5 x 11 in.)

- (Optional) Choose the orientation of the sheet. You can choose between **Portrait** or **Landscape**. The default option is portrait.

Before you can create Amazon Quick Sight pixel perfect reports, first get the pixel perfect reporting add-on for your Quick account. For more information on getting the pixel perfect reporting add-on, see [Get the Quick pixel perfect reports add-on](#).

For more information on formatting pixel perfect reports, see [Working with pixel perfect reports in Amazon Quick Sight](#).

6. Choose **Add**.

7. Create a visual. For more information about creating visuals, see [Adding visuals to Quick Sight analyses](#).

After you are done creating the analysis, you can iterate on it by modifying the visual, adding more visuals, adding scenes to the default story, or adding more stories.

You can also generate a complete multi-sheet analysis from a natural language prompt. For more information, see [Generating an analysis with natural language prompts](#).

Generating an analysis with natural language prompts

With Quick Sight, you can use artificial intelligence (AI) to author analyses in three ways within the analysis authoring experience. You can generate a full, multi-sheet analysis from a natural language prompt (Generate Analysis). You can generate a single sheet within an open analysis (Generate Sheet). You can also generate an analysis from an image. This topic describes the full-analysis flow. The following sections describe the two additional flows.

To generate a full analysis, describe the analysis or dashboard that you want, and select up to three datasets. Before generation begins, you can review and modify an interactive plan that outlines the proposed structure. When you generate, Quick Sight creates multiple organized sheets with visuals, filter controls, and calculated fields such as year-over-year growth and month-over-month comparisons.

The generated output is a native Quick Sight analysis. It works with existing publishing workflows, embedding patterns, CI/CD pipelines, and point-and-click editing in the analysis surface. After generation, you can refine each visual.

Prerequisites

To generate an analysis from a natural language prompt, you need the following:


- An AWS account
- Amazon Quick Enterprise Edition with at least one Author Pro user
- At least one dataset in your Quick Sight account

Generating an analysis


Use the following procedure to generate an analysis from a natural language prompt.


To generate an analysis from a natural language prompt

1. Do one of the following:
 - Open a dataset and choose **Generate analysis**.
 - From the **Analyses** page, choose **Generate analysis**.



2. Choose **Add data** to select one to three datasets for the analysis. If your data spans multiple tables (for example, orders in one dataset and products in another), you can select them together.







3. Enter a natural language prompt that describes the analysis that you want to create. You can describe the business questions that you want answered, the metrics that you care about, and how you want the information organized across sheets.

Example prompt:

"Create an operations dashboard showing order volume trends, revenue KPIs, delivery performance comparing estimated vs actual delivery dates, and product category breakdown by revenue and order count. Include calculated fields for total revenue, average order value, and month-over-month order growth."




4. Do one of the following:
 - Choose **Generate analysis** to begin generation immediately.
 - Choose **Preview analysis outline** to review an outline first.
5. Wait while Quick Sight analyzes your dataset structure and column statistics. Real-time progress updates display the current status.



Note


If you navigate away from the progress screen, you can check the generation status on the **Analyses** page by choosing the **Generations** tab. Choose the generation name to return to the progress screen.




6. Quick Sight presents a two-pane view:


- The left pane shows your initial prompt and a summary of the selected datasets.
- The right pane shows the proposed filter controls, sheets, and visuals planned for each sheet.

You can edit sheet names, add or remove visuals, adjust the plan, and refine the prompt before generating.



7. Choose **Generate**. Real-time progress updates display the current status. Generation takes 2 to 5 minutes depending on the number of sheets and visuals.






Publishing a generated analysis

After you are satisfied with the generated analysis, choose **Publish** to create a dashboard.

You can share the dashboard with other users, embed it in applications, or schedule email deliveries. For more information about publishing and sharing, see [Publishing dashboards](#) and [Sharing Quick Sight analyses](#).



Generating a single sheet

Within an open analysis, you can generate a single sheet with AI instead of generating a full analysis. Use the following procedure to generate a single sheet.

To generate a single sheet

1. In an open analysis, choose the + (plus) next to the sheet tabs.
2. In the **New sheet** dialog, select **Interactive sheet**, then choose **Generate**. Choosing **Add** creates a blank sheet to build by hand, and **Generate** starts the AI flow. **Generate** is available for interactive sheets only.
3. On the **Generate sheet** screen, select one or more datasets that are already part of the analysis, write a prompt that describes the sheet, and choose **Generate**.
4. Quick Sight adds the sheet with visuals selected for the data, filter controls, and calculated fields such as year-over-year growth and month-over-month comparisons.

Adding an image to your prompt

You can generate an analysis from an image of a dashboard, including dashboards from other business intelligence (BI) tools. Use the following procedure to add an image to your prompt.

To add an image to your prompt

1. From the **Visualize your data with AI** prompt, choose **Add image** (located next to **+ Add data**) to attach an image of a dashboard.
2. Quick Sight reads the image and builds the sheets and visuals that Quick Sight supports.

Adding a title and description to an analysis

In addition to the analysis name, you can add a title and description to an analysis. A useful title and description provides context about the information in the analysis.

Add a title and description

Use the following procedure to add a title and description to an analysis. Titles and descriptions can contain up to 1,024 characters. Titles and descriptions are not supported for pixel perfect reports.

To add a title and description to an analysis

1. On the analysis page, choose **Sheets** in the application bar and then choose **Add title**.
2. For **Sheet title**, enter a title and press **Enter**. To remove a title, choose **Sheets** in the application bar and then choose **Delete title**. Or, to remove the title, you can select the title and then choose the x-shaped delete icon.

To create a dynamic sheet title, you can add existing parameters to the sheet title. For more information, see [Using parameters in titles and descriptions in Amazon Quick](#).

3. Choose **Sheets** in the application bar, and then choose **Add description**.
4. In the description space that appears on the sheet, enter the description that you want and press **Enter**. To remove a description, choose **Sheets** in the application bar and then choose **Delete description**. Or, to remove the description, you can select the description and then choose the x-shaped delete icon.

Sharing Quick Sight analyses

You can share an analysis with one or more other users by emailing them a link, making it easy to collaborate and disseminate findings. You can only share an analysis with other users in your Quick account.

After you share an analysis, you can review the other users who have access to it, and also revoke access from any user.

Topics

- [Sharing an analysis](#)
- [Viewing the users that an analysis is shared with](#)
- [Revoking access to an analysis](#)

Sharing an analysis

Use the following procedure to share an analysis.

To share an analysis

1. Open the [Quick console](#).
2. Open the analysis that you want to change.
3. On the analysis page, choose **File** on the application bar, and then choose **Share**.

You can only share analyses with users or groups who are in your Quick account.

4. Add a user or group to share with. To do this, for **Type a user name or email**, enter the first user or group that you want to share this analysis with. Then choose **Share**. Repeat this step until you have entered information for everyone you want to share the analysis with.

To edit sharing permissions for this analysis, choose **Manage analysis permissions**.

The **Manage analysis permissions** screen appears. On this screen, choose **Invite user** to edit permissions and add more users or groups.

5. For **Permission**, choose the role to assign to each user or group. The role determines the permission level to grant to that user or group.
6. Choose **Share**.

The users that you have shared the analysis with get emails with a link to the analysis. Groups don't receive invitation emails.

Viewing the users that an analysis is shared with

If you have shared an analysis, you can use the following procedure to see which users or groups have access to it.

To view which users or groups have access to an analysis

1. Open the [Quick console](#).
2. Open the analysis that you want to change.
3. On the analysis page, choose **File** on the application bar, and then choose **Share**.
4. Choose **Manage analysis permissions**.
5. Review who this analysis has been shared with. You can search to locate a specific account by typing a search term. The search returns any user, group, or email address that contains the search term. Searching is case-sensitive, and wildcards are not supported. Delete the search term to view all users and groups.

Revoking access to an analysis

Use the following procedure to revoke access to an analysis.

To revoke access to an analysis

1. Open the [Quick console](#).
2. Open the analysis that you want to change.
3. On the analysis page, choose **File** on the application bar, and then choose **Share**.
4. Choose **Manage analysis permissions**.
5. Locate the user or group whose access you want to revoke, and then choose the trash-can icon next to the user or group.
6. Choose **Confirm**.

Renaming an analysis

Use the following procedure to rename an analysis.

To rename an analysis

1. Open the analysis that you want to rename.
2. In the **Analysis name** box in the application bar, select the current name and then enter a new name.

Duplicating analyses

You can duplicate analyses in Quick Sight. Use the following procedure to learn how.

To duplicate an analysis

1. From the Quick homepage, choose **Analyses**, and then open the analysis that you want to duplicate.
2. In the analysis, choose **Save as** in the application bar at upper right.
3. In the **Save a copy** page that opens, enter a name for the analysis, and then choose **Save**.

The new analysis opens. You can find the original analysis by returning to the Quick start page and selecting **Analyses**.

Customize date and time values of an analysis

In Amazon Quick, authors can set custom time zones and week start days of an analysis. When you set a custom week start or time zone, all visuals in the analysis that use datetime data are formatted to reflect the time zone or week start that the analysis uses.

Setting custom time zones in an analysis

Quick authors can use *custom time zones* to help manage data across multiple geographic regions. When you set a custom time zone, all visible dimensions, measures, calculated fields, and filters are converted to the chosen time zone at query run time. Daylight Savings Time (DST) adjustments are applied automatically to eliminate the need for time consuming workarounds that don't accurately handle historical dates.

Custom time zones refer to the use of IANA time zone abbreviations that represent specific geographic regions around the world. Each time zone is defined as an offset from Coordinated Universal Time (UTC). Time zones are different from simple offsets because they incorporate DST.

The default time zone for all analyses is UTC.

The following rules apply to time zones.

- **Datetime displays with a granularity that is lower than hour are converted to the selected time zone.** For example, if you set the timezone of an analysis to `America/New_York` (UTC-04:00), the datetime value `Dec.1, 2020 12:00am` in UTC+00:00 is converted and displayed as `Nov.30, 2020 7:00pm`. Daylight Savings Time (DST) is incorporated into the datetime conversion.
- **Datetime literals that are added to calculations or selected in filters honor the selected time zone of the analysis.** For example, if you manually enter a literal into a calculated field such as `01-01-2022 7:00pm`, or select a fixed filter time, Quick Sight applies the chosen timezone to the literal value.
- **Measures that are aggregated above the hour/minute granularity are aggregated based on the timezone that the analysis is set to.** When Quick Sight processes a dataset, all timestamps are initially converted at the lowest granularity level. Values are then aggregated based on the boundary of the selected time zone for the analysis. For example, a sum of hourly revenues at the day level with a UTC+00:00 time zone aggregate all hourly revenues from 12am-11pm for the UTC time zone. When you convert UTC+00:00 to `New_York` (UTC-04:00), all revenue datapoints are aggregated from 8:00pm-7:00pm(+1day) in UTC to correspond with the start and end of the day in `New_York` (UTC-04:00).
- **The `now()` function, rolling date filter, and parameters are converted to the chosen time zone.** Relative date filters, rolling date filters, and relative date parameters that use the `now()` function also honor the chosen time zone when they are applied to the visual. For example, when you select a relative filter such as `last week` or a rolling date filter such as `start of the month`, the chosen timezone is automatically applied to the filter to display the values `last week of New_York time zone` and `start of the month of New_York time zone`, respectively.

To set the custom time zone of an analysis

1. From the analysis that you want to change, navigate to the top menu and choose **Edit**.
2. Choose **Analysis settings**, and then choose **Date and time**.
3. Toggle **Convert time zone** ON, and choose the **Time zone** that you want.
4. Choose **Apply**.

When an analysis is assigned a time zone, an icon appears at the top of the analysis that indicates which time zone the analysis uses. This icon also appears on any dashboard that is published from the analysis.

Considerations

The following considerations apply to custom time zones.

- To use custom time zones, all datetime columns in a dataset must be normalized to UTC. If your datetime columns aren't normalized in your data source, you need to convert the columns in your data source before you can use this feature.
- For analyses that are not assigned a custom time zone, author and reader experiences are unaffected.
- Once a time zone is added to an analysis, the time zone is applied to all visuals and sheets in the analysis.
- Quick authors can choose only one time zone for an analysis. All dashboards that are published from the analysis use the time zone that the analysis uses. To create a dashboard that uses a different time zone than the one that the analysis uses, change the time zone of the analysis and republish the dashboard.
- Quick readers can't change the time zone of a dashboard.
- If you set the time zone of an analysis that uses a dataset that is stored in Direct Query and experience slow load times, consider storing the dataset in SPICE. SPICE is engineered to handle time zone conversions in a performant way.
- Custom time zones do not support the following database engines:
 - Timestream
 - OpenSearch Service
 - Teradata
 - SqlServer

Setting custom week start days in an analysis

Quick authors can define the week start day of an analysis to align their data with the schedule that their company or industry follows. When you set a custom week start day, all dimensions, calculated fields, and filters that are aggregated at the week level are calculated to align with the new week start day. The default week start day is Sunday.

To set the custom week start day of an analysis

1. From the analysis that you want to change, navigate to the top menu and choose **Edit**.
2. Choose **Analysis settings**, and then choose **Date and time**.
3. For **Custom start day**, choose the start day that you want.
4. Choose **Apply**.

Considerations

The following considerations apply to custom week start days.

- Datetime fields are converted at run time. When you work with calculated fields that use datetime values, define the fields at the analysis level instead of the dataset level.
- Once you choose a new week start day, the change is applied to all visuals and sheets in the analysis.
- Quick authors can choose only one week start day for an analysis. All dashboards that are published from the analysis use the week start day that the analysis uses. To create a dashboard that uses a different week start day than the one that the analysis uses, change the week start day of the analysis and republish the dashboard.
- Quick readers can't change the week start day of a dashboard.

The analysis menu

While working on an analysis, Amazon Quick provides several menu options. You use these menu options to efficiently perform tasks without needing to manually navigate through your analysis to find the assets that you want to change.

You can use these options to perform the following tasks.

- *File* – Perform analysis management tasks, including creating, sharing, and publishing. Authors can use this option to make changes across all sheets or visuals in an analysis.
- *Edit* – Navigate between changes that you make to the analysis. You can undo or redo changes that you make.
- *Data* – Manage datasets, data fields, and parameters. Changes that you make by using this option are applied to all sheets in the analysis.

- *Insert* – Use an ingress point where you can add visuals, text boxes, insights, reporting objects, filters, and parameters to an analysis. The content that you insert can be data or objects.
- *Sheets* – Manage the sheet settings of the analysis, including layout settings, actions to add or remove assets from a sheet, and sheet properties.
- *Objects* – Manage objects and their features, including style, canvas placement, sizing, card background, and borders. You can also manage these objects by using the **Properties** pane when working on a visual object.
- *Search* – Access the *Quick search* bar. Quick search is a search bar that will begin showing results for the asset you are searching for as you type. The suggested results continue to modify as you type until you see the result that you're looking for.

To use quick search, open the **Search** menu, and in the **Search analysis actions** box, begin typing a name or phrase associated with the asset you are trying to find.

Configure analysis settings

Amazon Quick authors can use the Analysis settings menu to configure the refresh and date time settings of an analysis. To access the Analysis Settings menu, choose **Edit**, and then choose **Analysis Settings**. The following settings can be configured in the Analysis settings menu:

Refresh settings

- **Reload visuals every time I switch sheets** – Use this setting to reload every visual in a Quick Sight analysis whenever the user switches to a different sheet in the analysis.
- **Update visuals manually** – Use this setting to only update applicable visuals in an analysis when the user applies their changes. When this setting is toggled on, the analysis loads the visuals blank by default because the queries won't be fired until the user selects the **UPDATE VISUALS** button located in the toolbar or on the impacted visuals. The **UPDATE VISUALS** button confirms that the user is finished with the filter and control choices that they want to apply to the affected visuals. The image below shows the **UPDATE VISUALS** button.

When **Update visuals manually** is toggled on, authors can still add visuals, edit visuals, and edit control selections, but the affected visuals won't update until the author applies the new changes. This allows authors to build analyses without increasing their database load and gives better control over which values are loaded in an analysis.

Date and time settings

- **Convert time zone** – Use this setting to convert all date field related visualizations, filters, and parameters to reflect the chosen time zone. All daylight savings adjustments are made automatically. For more information about time zone configuration, see [Customize date and time values of an analysis](#).
- **Start of the week** – Use this setting to choose the week start day for an analysis.

Interactivity

- Use this setting to highlight specific data points across visuals in a sheet. When you select or hover over a data point on a visual, related data across other visuals will stand out, while unrelated data is dimmed. Highlighting allows you to understand correlations, spot patterns, trends, and outliers, and facilitate stronger, more informed analyses. Select either **On selection** or **On hover** to turn highlighting on, or **No highlight** to turn it off.
- To customize highlighting on a per-sheet level see [Adding and managing sheets](#).

No data message customization

- Use this setting to control the message displayed when a visual in the analysis returns no data. Customizing this can give readers of the dashboards reasons why no data is shown and potential fixes. Parameters are supported for all text input.
- **Title and subtitle** – Use these settings to customize the text shown on the message title or subtitle and to toggle visibility of either option.
- **Hyperlink** – Use this setting to control the hyperlink displayed at the bottom of the message. The link label field controls the displayed text of the hyperlink, while the link URL field controls where the user is redirected when clicking the hyperlink text. Only http://, https://, and mailto: URL schemas are supported.
- **Pixel-perfect reporting** – In pixel-perfect (paginated) report output, hyperlinks are rendered as clickable links, not as static display text. To show the full URL in a printed report, enter the full URL as the link label so it remains visible in the static output.
- **Dataset or data refresh errors** – The no data message customization applies only to visuals that return zero rows of data. Dataset or data refresh errors are a separate error type and are not customizable as part of this feature.

Item limits for Amazon Quick Sight analyses in the Quick Sight APIs

Use the following table to review the current limits or quotas for different analysis items in Amazon Quick Sight that are created and managed with the Amazon Quick Sight APIs. If your analysis contains more than the supported number of analysis items, remove items to optimize the performance of the analysis. New analysis items cannot be added to an analysis that contains more than the supported number of analysis items.

Analysis item	Limit
Sheets	20 sheets per analysis
Visuals	50 visuals per sheet
Calculated fields	500 per analysis and 200 per dataset*
Bookmarks	200 per dashboard
Custom actions	10 per visual
Filter groups	2000 per analysis
Filters	20 filters per filter group
Parameters	400 per analysis
Controls	200 per sheet
Text boxes	100 per sheet
Image components	10 per sheet
Layer map visuals	5 per sheet

* The per dataset limit applies to calculations that were created in the analysis. Dataset level calculations are not included in this limit. For more information about dataset level calculations, see [Adding calculated fields](#).

Saving changes to analyses

When working on an analysis, you can set Autosave either on (the default) or off. When Autosave is on, your changes are automatically saved every minute or so. When Autosave is off, your changes aren't automatically saved, which means that you can make changes and pursue different lines of inquiry without permanently altering the analysis. If you decide that you want to save your results after all, re-enable Autosave. Your changes up to that point are then saved.

In either Autosave mode, you can undo or redo up to 200 changes that you make by choosing **Undo** or **Redo** on the application bar.

Changing the Autosave mode

Changed to To change the Autosave mode for an analysis, choose **File** and then choose **Autosave On** or **Autosave OFF**.

When Autosave can't save changes

Suppose that one of the following things occurs:

- Autosave is on and another user makes a conflicting change to the analysis.
- Autosave is on and there is a service failure, such that your most recent changes can't be saved.
- Autosave is off, you turn it on, and one of the backlogged changes now being saved to the server conflicts with another user's changes.

In this case, Amazon Quick Sight gives you the option to do one of two things. You can either let Amazon Quick Sight turn Autosave off and continue working in unsaved mode, or reload the analysis from the server and then redo your most recent changes.

If your client authentication expires while you are editing an analysis, you are directed to sign in again. On successful sign-in, you are directed back to the analysis where you can continue working normally.

If your permissions on the analysis are revoked while you are editing it, you can't make any further changes.

Exporting data from Quick Sight analyses

Note

Export files can directly return information from the dataset import. This makes the files vulnerable to CSV injection if the imported data contains formulas or commands. For this reason, export files can prompt security warnings. To avoid malicious activity, turn off links and macros when reading exported files.

You can export data from an analysis to a CSV or PDF file. To export data from an analysis or dashboard to a CSV file, follow the procedure in [Exporting data from visuals](#).

Use the procedure below to export an analysis as a PDF.

1. From the analysis that you want to export, choose **File > Export to PDF**. Quick Sight begins to prepare the analysis for download.
2. Choose **VIEW EXPORTS** in the blue pop-up to open the **Exports** pane on the right.
3. Choose **DOWNLOAD** in the green pop-up.
4. To see all analyses or reports that are ready to download, choose **File** then **Exports**. The Exports panel will open on the right side of the screen. Select **Click to download** next to the file that you want to save to your preferred location.

The process for exporting to a PDF works the same way for both dashboards and analyses.

You can also attach a PDF to dashboard email reports. For more information, see [Scheduling and sending Quick Sight reports by email](#).

Deleting an analysis

If you have the permissions to do so, you can delete an analysis from the **Analyses** page. When you delete an analysis, it doesn't affect any dashboards that are based on that analysis. They continue to show the deleted analysis, but you can't make changes to the analysis once you delete it. Navigate to the Analyses page and find the analysis that you want to remove. Choose the details icon (:) on the analysis, then choose **Delete**. Confirm your choice by choosing **Delete** again. You can't undo this action.

Adding and managing sheets

A *sheet* is a set of visuals that are viewed together in a single page. When you create an analysis, you place visuals in the workspace on a sheet. You can imagine this as a sheet from a newspaper, except that it is filled with data visualizations. You can add more sheets, and make them work separately or together in your analysis.

The top sheet, also called the default sheet, is the one on the far left. This sheet displays on top in an analysis or dashboard. Each analysis can contain up to 20 sheets.

You can share analyses and publish dashboards with multiple sheets. You can also schedule email reports for any combination of sheets in an analysis.

When you create a new analysis or a new sheet in an existing analysis, you choose whether to make the new sheet an **Interactive sheet** or a **Pixel perfect report**. This way, you can have analyses for interactive sheets only, analyses for pixel perfect reports only, or you can have an analysis that includes both interactive sheets and pixel perfect reports.

An *Interactive Sheet* is a collection of data expressed in visuals that users can interact with when the sheet is published to a dashboard. Amazon Quick authors can add different controls and filters to their interactive sheets. Dashboard viewers can use these to gain detailed information from the published data. For more information on interactive sheets, see [Working with interactive sheets in Amazon Quick Sight](#).

A *Paginated Report* is a collection of tables, charts, and visuals that are used to convey business critical information, such as daily transaction summaries or weekly business reports. In order to create pixel perfect reports in Quick Sight, add the **Pixel perfect reporting Add-on** to your Quick account. To get the **Pixel perfect reporting Add-on** and start working with pixel perfect reports, see [Working with pixel perfect reports in Amazon Quick Sight](#).

Use the following list of actions to work with sheets:

- To add a new sheet, choose the plus sign (+) to the right of the sheet tabs, choose the type of sheet that you want, and then choose **ADD**.
- To rename a sheet, choose the name of the sheet and start typing. **Rename** is also available from the sheet menu.
- To duplicate a sheet, choose the name of the sheet, then choose **Duplicate** from the sheet menu. You can only duplicate a sheet if **Autosave** is turned on.

- To duplicate an interactive sheet and convert it to a pixel perfect report, choose the name of the sheet, then choose **Duplicate to report** from the sheet menu. You can't convert a pixel perfect report to an interactive sheet.
- To delete a sheet, choose the name of the sheet, then choose **Delete** from the sheet menu. You can't delete the sheet if it's the only sheet in the analysis.
- To change the order of the sheets, choose the name of the sheet and drag it to a new position.
- To copy a visual to a new sheet, choose **Duplicate visual to** from the on-visual menu. Then choose the target sheet. Filters exist only on the sheet that you create them on. To duplicate filters, recreate them on the target sheet.
- To highlight specific data points across visuals in a sheet, go to the **Sheets** tab and select **Layout Settings**. Under the **Interactivity** section, select either **On selection** or **On hover** to turn highlighting on, or **No highlight** to turn it off. By default, sheet highlighting follows the same settings as analysis highlighting.

When you select or hover over a data point on a visual, related data across other visuals will stand out, while unrelated data is dimmed. Highlighting allows you to understand correlations, spot patterns, trends, and outliers, and facilitate stronger, more informed analyses.

You can use the parameter controls on the top sheet to control multiple sheets. To do this, open each sheet that you want to work with the parameter. Then add a filter that uses the same parameter used in the control on the top sheet. Or, if you want a new sheet to operate independently, you can add parameters and parameter controls to it that are separate from those on the top sheet.

Working with interactive sheets in Amazon Quick Sight

An *Interactive Sheet* is a collection of data expressed in visuals that users can interact with when the sheet is published to a dashboard. Amazon Quick authors can add different layouts, controls, and filters to their interactive sheets that dashboard viewers can use to gain detailed information from the published data. By default, every sheet in an analysis is an interactive sheet. If your account doesn't have the **Pixel perfect reporting Add-on**, you can only create and publish interactive sheets.

For more information on creating an interactive sheet, see [Starting an analysis in Quick Sight](#).

For more information on formatting interactive sheets, see the following topics.

Topics

- [Customizing dashboard layouts in Amazon Quick Sight](#)
- [Parameters in Amazon Quick](#)
- [Using custom actions for filtering and navigating](#)

Customizing dashboard layouts in Amazon Quick Sight

You can customize a dashboard's layout to organize your data to fit your business requirements. You can choose from three dashboard layouts. You can also change the size, background color, border color, and interactions of a visual to create a fully customized dashboard.

Use the following topics to learn more about customizing dashboards and visuals.

Topics


- [Types of layout](#)
- [Choosing a layout](#)
- [Customizing visuals in a free-form layout](#)
- [Conditional rules](#)

Types of layout

There are three dashboard layout designs you can choose from: **Tiled**, **Free-form**, and **Classic**.

Tiled layout

Visuals in a **Tiled** layout snap to a grid with standard spacing and alignment. You can make visuals any size and place them wherever you want within a dashboard, but visuals can't overlap.




Dashboards are displayed as designed, with options to fit to screen or view at actual size. You can also fit an entire dashboard to your window by choosing **Fit to window** for **View**. This option was previously called **Optimized**.

Note

On mobile devices, tiled layout dashboards appear as a single column in portrait mode or exactly as designed in landscape mode.

Free-form layout

Visuals in a **Free-form** layout can be placed anywhere in your dashboard using precise coordinates. You can drag a visual to the exact place you want, or you can enter the coordinates of the visual's location. Use the following procedure to enter the exact coordinates of the visual's location.



Dashboards are displayed the way that you choose to design them, with options to fit to screen or to view at its actual size. You can optimize free-form layouts for viewing at specific resolutions, with the default being 1,600 pixels. You can also fit an entire dashboard to a browser's window by choosing **Fit to window** for **View**.

Note

Dashboards with optimized resolutions might appear bigger or smaller on a viewer's computer if the viewer's computer resolution doesn't equal the set resolution of the dashboard.

Switching from **Free-form** to another layout might cause some visual elements to shift. On mobile devices, **Free-form** dashboards appear as published with no changes to the layout.

Classic layout

Visuals in a **Classic** layout snap to a grid with standard spacing and alignment. Dashboards hide data or change formatting to fit smaller screen sizes. For example, if you change a visual to make it considerably smaller, the on-visual menu and editors are hidden so that the chart elements have more room to display. Bar chart visuals can also display fewer data points.

If you reduce the size of the browser window, Amazon Quick Sight resizes and if necessary reorders visuals for optimal display. For example, smaller visuals that were side by side might be displayed sequentially. The original layout is restored when the size of the browser window is increased again.

Note

On mobile devices, classic layout dashboards appear as a single column or exactly as designed in landscape mode.

Choosing a layout

To change a dashboard's layout

1. Open the [Quick console](#).
2. From the Quick homepage, choose **Analyses**, and then choose the analysis that you want to change.
3. On the analysis page, choose **Edit** and then choose **Analysis Settings**.
4. Expand **Sheet layout** and choose the layout that you want to use.
5. When finished, choose **Apply**.

Customizing visuals in a free-form layout

You can use the free-form layout to fully customize the color, size, location, and visibility of each visual in a dashboard.

Organizing visuals

Besides dragging a visual to its preferred location within a dashboard, there are many different ways to move a visual to the exact location it needs to be.

To enter the coordinates of the visual's location

1. Choose the visual that you want.
2. On the menu in the upper-right corner of the visual, select the **Format visual** icon.
3. In the **Properties** pane that opens, choose **Placement**.
4. Enter the **X** and **Y** coordinates of the location you want to place your visual. You can also adjust the size of the visual by entering **Width** and **Height** values.

Selected visuals can also be moved pixel-by-pixel using your keyboard's arrow keys.

You can overlay visuals on top of one another to create multi-layered visuals that show data.

Visuals can be organized into multiple layers that can be manually moved to the front and back.

To move overlaid visuals to the front and back

1. Choose the visual that you want.
2. On the three-dot menu in the upper-right hand side of the visual, choose **Menu options**.
3. For **Menu options**, choose from the following:
 - **Send to back** sends the visual to the back.
 - **Send backward** sends the visual one layer back.
 - **Bring forward** rings the visual one layer forward.
 - **Bring to front** brings a visual to the front.

Changing a visual's background color

The colors of a visual's background, border, and selection frame can be customized in the **Display settings** pane of the **Properties** pane.

To change the color of a visual's background, border, or selection frame

1. Choose the visual that you want to change.
2. On the menu in the upper-right hand side of the visual, choose the **Properties** icon.
3. In the **Properties** pane that appears on the left, choose **Display settings**.
4. Navigate to the **Card style** section and perform one or more of the available actions:

- To change the color of a visual's background, choose the **Background** color box, and then choose the color that you want.
- To change the color of a visual's border, choose the **Border** color box, and then choose the color that you want.
- To change the color of a visual's selection frame, choose the **Selection** color box, and then choose the color that you want.

If you want to use a custom color for your visual's background, border, or selection frame, choose the color box of the property that you want to change, and then choose **Custom color**. In the **Custom color** window that appears, choose your custom color or enter the color's hexadecimal code. When you are finished, choose **Apply**.

You can also reset a visual's customized background back to its default appearance.

To reset the appearance of a visual

1. Choose the visual that you want to change.
2. On the menu in the upper-right hand side of the visual, choose the **Properties** icon.
3. In the **Properties** pane that appears on the left, choose **Display settings**.
4. Choose the color that you want to reset, and then choose **Reset to default**.

Hiding visual backgrounds, borders, and selection colors

You can also choose not to show the background border, or selection color of a visual. This is useful for when you want to overlap multiple visuals. You can hide a visual's background, border, and selection colors by choosing the eye icon next to the **Border**, **Background**, or **Selection** color boxes. You can also remove a visual's loading animation by clearing the **Show loading animation** box.

Disabling visual menus

Use the **Interactions** panel of the **Properties** pane to hide the **Context** menu and **On-visual** menu from selected visuals. You can hide secondary visual menus to make the visual less crowded or to make a visual act like an overlay.

The **Context** menu opens on data-point clicks. Common actions in the **Context menu** include **Focus**, **Exclude**, and **Drill-down**.

The **On-visual** menu appears on the top-right side of a visual. The **On-visual** menu is used to access the **Properties** pane, **Maximize** the visual, access the **Menu options** panel, and review an **Anomaly insight**.

You can turn off the secondary visual menus by clearing the **Context menu** and **On-visual menu** options.

Note

You can't preview changes to the **Interactions** panel in **Analyses**. Publish the dashboard to view your changes.

Conditional rules

This feature is currently available with the **Free-form** layout. Conditional rules are used to hide or show visuals when specific conditions are met. This can be useful when you have multiple versions of the same visual overlapped with each other and want the dashboard viewer to see a version that best represents the parameter value they select.

Conditional rules use parameters and parameter controls to hide and show visuals. Parameters are named variables that can transfer a value for use by an action or an object. This feature supports string and number parameters. To make the parameters accessible to the dashboard viewer, you add a parameter control. A parameter control allows users to choose a value to use in a predefined filter or URL action. For more information about parameters and parameter controls, see [Parameters in Amazon Quick](#).

Use the sections below to set up and use conditional rules.

Topics

- [Hiding a visual by default](#)
- [Setting a conditional rule](#)
- [Using conditional rules](#)

Hiding a visual by default

In the **Interactions** pane of the **Properties** pane, you can choose to hide a visual by default. Doing this can be useful if you want the viewer to only see visuals based on specific conditions.

To hide a visual by default

1. From the Quick homepage, choose **Analyses**, and then choose the analysis that you want to customize.
2. Choose the visual that you want to add a rule to.
3. On the menu in the upper-right hand side of the visual, choose **Properties**.
4. In the **Properties** pane that opens, choose **Interactions** and open the **Rules** dropdown.
5. In the **Rules** menu, choose **Hide this visual by default**.

Hidden visuals appear fully hidden in a viewing dashboard. In the **Analyses** pane, hidden visuals are visible with the message "Hidden based on rule". With this display, you can see where all of a dashboard's visuals are located.

Note

You can't create conditional rules that hide visuals that are already hidden by default or that show visuals that already appear by default. If you change the default appearance of a visual, existing rules that contradict the new default appearance will be disabled.

Setting a conditional rule

When you set up a conditional rule, you create a conditional statement that will hide or show a visual when a specific condition is met. You can currently create conditional rules that hide or show a visual. If you want to create a conditional rule that makes a hidden visual appear, choose **Hide this visual by default** in the **Rules** menu of the **Properties** pane.

Note

Before you begin, make a parameter and a corresponding parameter control to base your new conditional rule on. Supported parameters are string parameters and number parameters. For more information about parameters and parameter controls, see [Parameters in Amazon Quick](#).

To set a conditional rule

1. From the Quick homepage, choose **Analyses**, and then choose the analysis you want to customize.
2. Choose the visual that you want to add a rule to.
3. On the menu in the upper-right hand side of the visual, choose **Properties**.
4. In the **Properties** pane that appears on the left, choose **Interactions**, and then choose **Rules**.
5. Choose **ADD RULE**.
6. In the first menu in the **Add rule** pane, choose the parameter you want.
7. In the second menu in the **Add rule** pane, choose which condition you want. For string parameters, supported conditions are **Equals**, **Starts with**, **Contains**, and **Does not equal**. For number parameters, supported conditions are **Equals**, **Starts with**, **Contains**, and **Does not equal**.
8. Enter the value you want the conditional rule to meet.

Note

Values are case-sensitive.

9. Choose **Add rule** to apply the new conditional rule to the visual. To cancel the rule, choose **Cancel**.

Conditional rules can also be edited and deleted.

To edit a conditional rule

1. On the menu in the upper-right hand side of the visual, choose **Properties**.
2. In the **Properties** pane that appears on the left, choose **Interactions**, and then choose **Rules**.
3. Choose the menu icon on the right-hand side of the rule you want to edit, and choose **Edit**.
4. Make the changes that you want and choose **Save**.

To delete a conditional rule

1. On the menu in the upper-right hand side of the visual, choose **Properties**.
2. In the **Properties** pane that appears on the left, choose **Interactions**, and then choose **Rules**.

3. Choose the menu icon on the right-hand side of the rule you want to edit and choose **Delete**.

Using conditional rules

Once you have set up a conditional rule that is connected to a parameter and a parameter control, you can use the parameter control to enable or disable the conditional rules you have set.

To enable a conditional rule

1. From the Quick homepage, choose **Analyses**, and then choose the analysis you want to customize.
2. On the **Controls** bar, choose the dropdown icon.
3. Choose the parameter control associated with the conditional rule you created.
4. Choose the value associated with the conditional rule that you created from the parameter's menu. You can also enter the value that you want into the **Search value** box.

Note

Values are case-sensitive.

Selecting the correct value causes the visual to appear or disappear depending on the rule you set.

You can also bring a parameter control to the sheet your visual is on. This is useful when you want a parameter control to be next to the visual it is associated with or when you want to add a conditional rule to the control so it appears only when specific conditions are met.

To bring a parameter control to a sheet

1. From the Quick homepage, choose **Analyses**, and then choose the analysis you want to customize.
2. On the **Controls** bar, choose the control that you want to move.
3. At the upper right-hand side of the control, open the **Menu options** menu.
4. Choose **Move to sheet**.

To move a parameter control back to the Controls bar

1. On your dashboard, select the parameter control you want to move.
2. On the upper right-hand side of the control, open the **Menu options** menu.
3. Choose **Move to top of sheet**.

Parameters in Amazon Quick

Parameters are named variables that can transfer a value for use by an action or an object. By using parameters, you can create an easier way for a dashboard user to interact with dashboard features in a less technical way. Parameters can also connect one dashboard to another, allowing a dashboard user to drill down into data that's in a different analysis.

For example, a dashboard user can use a list to choose a value. That value sets a parameter that in turn sets a filter, calculation, or URL action to the chosen value. Then the visuals in the dashboard react to the user's choices.

To make the parameters accessible to the dashboard viewer, you add a parameter control. You can set up cascading controls, so that a selection in one control filters the options that display in another control. A control can appear as a list of options, a slider, or a text entry area. If you don't create a control, you can still pass a value to your parameter in the dashboard URL.

For a parameter to work, it needs to be connected to something in your analysis, regardless of whether it has a related control. You can reference parameters in the following:

- Calculated fields (except for multivalue parameters)
- Filters
- Dashboard and analysis URLs
- Actions
- Titles and descriptions throughout an analysis

Some ways that you can use parameters are the following:

- Using a calculation, you can transform data that is shown in an analysis.
- If you add a control with a filter to an analysis you are publishing, the dashboard users can filter the data without creating their own filters.
- Using controls and custom actions, you can let dashboard users set values for the URL actions.

Topics

- [Setting up parameters in Amazon Quick](#)
- [Using a control with a parameter in Amazon Quick](#)
- [Creating parameter defaults in Amazon Quick](#)
- [Connecting to parameters in Amazon Quick](#)

Setting up parameters in Amazon Quick

Use the following procedure to create or edit a basic parameter.

To create or edit a basic parameter

1. Choose an analysis to work with, then decide which field you want to parameterize.
2. Choose the **Parameters** icon from the icon list at the top of the page.
3. Add a new parameter by choosing the plus sign (+ **Add**) near the top of the pane.

Edit an existing parameter by first choosing the v-shaped icon near the parameter name and then choosing **Edit parameter**.

4. For **Name**, enter an alphanumeric value for the parameter.
5. For **Data type**, choose **String**, **Number**, **Integer**, or **Datetime**, and then complete the following steps.

- If you choose **String**, **Number**, or **Integer**, do the following:

1. For **Values**, choose **Single value** or **Multiple values**.

Choose the single value option for parameters that can contain only one value. Choose multiple values for parameters that can contain one or more values. Multivalue parameters can't be datetime data types. They also don't support dynamic default values.

To switch an existing parameter between single and multiple values, delete and recreate the parameter.

2. (Optional) For **Static default value** or **Static multiple default values**, enter one or more values.

This type of static value is used during the first page load if a dynamic default value or URL parameter isn't provided.

3. (Optional) Choose **Show as blank by default**.

Select this option to show the default value for multivalue lists as blank. This option only applies to multivalue parameters.

- If you choose **Datetime**, do the following:

1. For **Time granularity**, choose **Day**, **Hour**, **Minute**, or **Second**.

2. For **Default date**, select either **Fixed date** or **Relative date**, and then do the following:

- If you select **Fixed date**, enter a date and time by using the date and time picker.
- If you select **Relative date**, choose a rolling date. You can choose **Today**, **Yesterday**, or you can specify the **Filter condition** (start of or end of), **Range** (this, previous, or next), and **Period** (year, quarter, month, week, or day).

6. (Optional) Choose **Set a dynamic default** to create a default that is user-specific.

A *dynamic default* is a per-user default value for the first page load of the dashboard. Use a dynamic default to create a personalized view for each user.

Calculated fields can't be used as dynamic defaults.

Dynamic defaults don't prevent a user from selecting a different value. If you want to secure the data, you can add row-level locking. For more information, see [Using row-level security with user-based rules to restrict access to a dataset](#).

This option only appears if you choose a single value parameter. Multivalue parameters can't have dynamic defaults.

Note

If you choose a multivalue parameter, the screen changes to remove the default options. Instead, you see a box with the text **Enter values you want to use for this control**. You can enter multiple values in this box, each on a single line. These values are used as the default selected values in the parameter control. The values here are unioned with what you choose to enter for the parameter control. For more information on parameter controls, see [Parameter Controls](#).

- ### 7. (Optional) Set a reserved value to determine the value of the **Select all** value. The *reserved value* of a parameter is the value that is assigned to a parameter when you choose **Select all** as its value. When you set up a specific reserved value for your parameter, that value is no longer

considered a valid parameter value in your dataset. The reserved value can't be used in any *parameter consumers*, such as filters, controls, and calculated fields, and custom actions. Also, it does not appear in the parameter control list. You can choose from **Recommended value**, **Null**, and **Custom value**. **Recommended value** is the default. If you choose **Recommended value**, the reserved value is set to the following values based on the value type:

- Strings: "ALL_VALUES"
- Numbers: "Long.MIN_VALUE"-9,223,372,036,854,775,808
- "Integers: Int.MIN_VALUE"-2147483648

To set a reserved value in your new parameter, choose the **Advanced settings** dropdown list in either the **Create a new parameter** page or the **Edit parameter** page and select the value that you want.

8. Choose **Create** or **Update** to complete creating or updating the parameter.

After you create a parameter, you can use it in a variety of ways. You can create a control (such as a button) so that you can choose a value for your parameter. For more information, see the following sections.

Using a control with a parameter in Amazon Quick

In dashboards, parameter controls appear at the top of the data sheet, which contains a set of visuals. Providing a control allows users to choose a value to use in a predefined filter or URL action. Dashboard users can use controls to apply filtering across all visuals datasets on a dashboard, without having to create the filters themselves.

The following rules apply:

- To create or edit a control for a parameter, make sure that the parameter exists.
- Multiselect list controls are compatible with analysis URLs, dashboard URLs, custom actions, and custom filters. The filter must be either equal or not equal to the values provided. No other comparisons are supported.
- Lists show up to 1,000 values. If there are more than 1,000 distinct values, a search box appears so you can filter the list. When the filtered list contains less than 1,001 values, the contents of the list appear as line items.

- The **Style** option displays only the style types that are appropriate for the parameter's data type and single or multivalue setting. If the style that you want to use isn't in the list, recreate your parameter with the appropriate settings and try again.
- If your parameter links to a dataset field, it must be an actual field. Calculated fields aren't supported.
- The values display alphabetically in the control, unless there are more than 1,000 distinct values. Then the control displays a search box instead. Each time you search for the value you want to use, it initiates a new query. If the results contain more than 1,000 values, you can scroll through the values with pagination. Wildcard search is supported. To learn more about wildcard search, see [Using wildcard search](#).

Use the following procedure to create or edit a control for an existing parameter.

To create or edit a control for an existing parameter

1. Choose an existing parameter's context menu, the  icon near the parameter name, and choose **Add control**.
2. Enter a name to give the new control a label. This label appears at the top of the workspace, and later at the top of the sheet that a dashboard displays on.
3. Choose a style for the control from the following:

- **Text field**

A text field lets you type in their own value. A text field works with numbers and text (strings).

- **Text field - multiline**

A multiline text field lets you type in their own values. With this option, you can choose to separate values you enter into the parameter control by a line break, comma, pipe (|), or semicolon. A text field works with numbers and text (strings).

- **Dropdown**

A dropdown list control that you can use to select a single value. A list control works with numbers and text (strings).

- **Dropdown multiselect**

A list control that you can use to select multiple values. A list control works with numbers and text (strings).

- **List**

A list control that you can use to select a single value. A list control works with numbers and text (strings).

- **List - multiselect**

A list control that you can use to select multiple values. A list control works with numbers and text (strings).

- **Slider**

A slider lets you select a numeric value by sliding the control from one end of the bar to another. A slider works with numbers.

- **Date-picker**

Using a date-picker, you can choose a date from a calendar control. When you choose to add a date-picker control, you can customize how to format dates in the control. To do so, for **Date format**, enter the date format that you want using the tokens described in [Customizing date formats in Quick](#).

4. (Optional) If you choose a dropdown control, the screen expands so you can choose the values to display. You can either specify a list of values, or use a field in a dataset. Choose one of the following:

- **Specific values**

To create a list of specific values, type in one per line, with no separating spaces or commas, as shown in the following screenshot.

In the control, the values display alphabetically, not in the order that you typed them.

- **Link to a data set field**

To link to a field, choose the dataset that contains your field, then choose the field from the list.

If you change the default values in the parameter, choose **Reset** on the control to show the new values.

The values that you choose here are unioned with the static default values in the parameter settings.

5. (Optional) Enable the option **Hide [ALL] option from the control if the parameter has a default configured**. Doing this shows only the data values and removes the option to select all items in the control. If you don't configure a static default on the parameter, this option doesn't work. You can add a default after adding a control by choosing the parameter, and selecting **Edit parameter**.
6. (Optional) You can limit the values displayed in the controls, so they only show values that are valid for what is selected in other controls. This is called a cascading control.

To create one, choose **Show relevant values only**. Choose one or more controls that can change what displays in this control.

When creating cascading controls, the following limitations apply.

- Cascading controls must be tied to dataset columns from the same dataset.
 - The child control must be a dropdown or list control.
 - For parameter controls, the child control must be linked to a dataset column.
 - For filter controls, the child control must be linked to a filter (instead of showing only specific values).
 - The parent control must be one of the following.
 - A string, integer, or numeric parameter control.
 - A string filter control (EXCLUDING Top-Bottom filters).
 - A non-aggregated numeric filter control.
 - A date filter control (EXCLUDING Top-Bottom filters).
7. When you finish choosing options for your control, choose **Add**.

The finished control appears at the top of the workspace. The context menu, shaped like a v, offers four options:

- **Reset** restores the user's selection to its default state.
- **Refresh list** applies only to drop-downs that are linked to a field in a dataset. Choosing **Refresh list** queries the data to check for changes. Data used in the control is cached.
- **Edit** reopens the control creation screen so that you can change your settings.

Once you have the **Edit control** pane open, you can click on different visuals and controls to view formatting data for the specific visual or control. For more information about formatting a visual, see [Formatting in Amazon Quick](#).

- **Delete** removes the control. You can recreate it by choosing the parameter context menu.

Configuring custom sort for controls

You can configure a custom sort order for parameter controls to change how values appear in the control list. By default, controls display values in user-defined order. You can customize this to sort in ascending or descending order or sort by a related field in the dataset.

To configure custom sort for a control

1. Choose the parameter control in the sheet.
2. Choose the pencil icon to open the **Format control** pane.
3. In the **Format control** pane, locate the **Sort** section.
4. Configure the sort options. Note that the sort options depend on whether the control uses specific values or is linked to a dataset field.

Sort options by control type

Sorting controls with specific values

If the control uses specific values, choose one of the following sort options:

- **Ascending (A–Z, 0–9)** – Sorts values in ascending order.
- **Descending (Z–A, 9–0)** – Sorts values in descending order.
- **User-defined order** – Displays the values in the exact order that you entered them, without any automatic sorting. This is the default.

Sorting controls linked to a dataset field

If the control is linked to a dataset field, you can sort by the control's own column or by another field in the same dataset.

- **Sort by control column** – Sorts values based on the field that the control is linked to. Choose a sort direction (**Ascending** or **Descending**) and optionally choose an aggregation function. Choose **No aggregation** to sort by the raw column values.
- **Sort by another field** – Sorts values based on a column in any analysis dataset. For calculated fields, only scalar (non-aggregated) calculated columns are available. Choose a sort-by field, a sort direction, and an aggregation function. This is useful for ordering values by a related metric, such as sorting regions by total revenue.

Note

The available aggregation functions depend on the data type of the selected column. For numeric fields, options include **No aggregation**, **Sum**, **Average**, **Count**, **Distinct count**, **Min**, **Max**, **Median**, **Percentile**, **Var**, and **Stdev**. For non-numeric fields, options include **No aggregation**, **Count**, and **Distinct count**.

In the workspace, you can also resize and rearrange your controls. The dashboard users see them as you do, except without being able to edit or delete them.

Creating parameter defaults in Amazon Quick

Use this section to learn more about the types of parameter defaults that are available, and how to set up each of them.

Each field can have a parameter and a control associated with it. When someone views a dashboard or email report, any sheet control that has a static default value configured uses the static default. The default value can change how data is filtered, how custom actions behave, and what text displays in a dynamic sheet title. Email reports also support dynamic defaults.

The simplest default is a static (unchanging) default, which shows the same value to everyone. As the designer of the dashboard, you choose the default value. It can't be changed by the person using the dashboard. However, that person can choose any value from the controls. Setting a default doesn't change this. To restrict the values that a person can select, consider using row-level security. For more information, see [Using row-level security with user-based rules to restrict access to a dataset](#).

To create or edit a static default value that applies to everyone's dashboard view

1. Choose the context menu (v) by the parameter that you want to edit, or create a new parameter by following the steps in [Setting up parameters in Amazon Quick](#).
2. Enter a value for **Static default value** to set a static default.

To display a different default depending on who is viewing the dashboard, you create a dynamic default parameter (DDP). Using dynamic defaults involves some preparation to map people to their assigned defaults. First, you need to create a database query or a data file that contains information about the people, the fields, and the default values to display. You add this to a dataset, then add the dataset to your analysis. Following, you can find procedures that you can use to gather information, create the dataset, and add the dynamic default to the parameter.

Use the following guidelines when creating a dataset for dynamic default values:

- We recommend that you use a single dataset to contain all dynamic default definitions for a logical grouping of users or groups. If you can, maintain them in a single table or file.
- We also recommend that the fields in your dataset have names that closely match the field names in the analysis. Not all dataset fields need to be part of the analysis, for example if you're using the same dataset for the defaults in multiple dashboards. The fields can be in any order.
- We don't recommend that you combine both user and group names in the same column or even in the same dataset. This kind of configuration is more work to maintain and troubleshoot.
- If you use a comma-delimited file to create your dataset, make sure to remove any space between values in the file. The following example shows the correct comma-separated value (CSV) format. Enclose text (strings) that include nonalphanumeric characters—like spaces, apostrophes, and so on—in single or double quotation marks. You can enclose fields that are dates or times in quotation marks, but it isn't required. You can enclose numeric fields in quotation marks, for example if the numbers contain special characters, as shown following.

```
"Value includes spaces","Field contains ' other characters",12345.6789,"20200808"
ValueWithoutSpaces,"1000,67","Value 3",2020-AUG-08
```

- After you create the dataset, make sure to double-check the data types that Quick selects for the fields.

Before you begin, you need a list of the user or group names for the people who are going to have dynamic defaults. To generate a list of users or groups, you can use the AWS CLI to get the

information. To run CLI commands, make sure that you have the AWS CLI installed and configured. For more information, see [Installing the AWS CLI](#) in the *AWS CLI User Guide*.

This is just one example of how to get a list of user or group names. Use whatever method works best for you.

To identify people for a dynamic default parameter (DDP)

- List either individual user names or group names:
 - To list individual user names, include a column that identifies the people for your DDP. This column should contain each person's system user name that they use to connect from your identity provider to Quick. This user name is often the same as a person's email alias before the @ sign, but not always.

To get a list of users, use the [ListUsers](#) Quick API operation or AWS CLI command. The CLI command is shown in the following example. Specify the AWS Region for your identity provider, for example `us-east-1`.

```
awsacct1="111111111111"
namespace="default"
region="us-east-1"

aws quicksight list-users --aws-account-id $awsacct1 --namespace $namespace --
region $region
```

The following example alters the previous command by adding a query that limits the results to active users.

```
awsacct1="111111111111"
namespace="default"
region="us-east-1"

aws quicksight list-users --aws-account-id $awsacct1 --namespace $namespace --
region $region --query 'UserList[?Active==`true`]'
```

The result set looks similar to the following sample. This example is an excerpt from JSON output (`--output json`). People who have federated user names have principal IDs that start with the word `federated`.

```
[
  {
    "Arn": "arn:aws:quicksight:us-east-1:111111111111:user/default/
anacasilva",
    "UserName": "anacarolinasilva",
    "Email": "anacasilva@example.com",
    "Role": "ADMIN",
    "Active": true,
    "PrincipalId": "federated/iam/AIDAJ64EIEIOPX5CEIEIO"
  },
  {
    "Arn": "arn:aws:quicksight:us-east-1:111111111111:user/default/Reader/
liujie-stargate",
    "UserName": "Reader/liujie-stargate",
    "Role": "READER",
    "Active": true,
    "PrincipalId": "federated/iam/AR0AIJSEIEIOMXTZEIEIO:liujie-stargate"
  },
  {
    "Arn": "arn:aws:quicksight:us-east-1:111111111111:user/default/embedding/
cxoportall",
    "UserName": "embedding/cxoportal",
    "Email": "saanvisarkar@example.com",
    "Role": "AUTHOR",
    "Active": true,
    "PrincipalId": "federated/iam/AR0AJTGEIEIOWB6BEIEIO:cxoportall"
  },
  {
    "Arn": "arn:aws:quicksight:us-east-1:111111111111:user/default/
zhangwei@example.com",
    "UserName": "zhangwei@example.com",
    "Email": "zhangwei@example.com",
    "Role": "AUTHOR",
    "Active": true,
    "PrincipalId": "user/d-96123-example-id-1123"
  }
]
```

- To list group names, include a column that identifies the groups containing the user names for your DDP. This column should contain the system group names that are used to connect from your identity provider to Quick. To identify groups that you can add to the dataset, use one or more of the following Quick API operations or CLI commands:

- [ListGroups](#) – Lists Quick groups by AWS account ID and namespace for the AWS Region that contains your identity provider.
- [ListGroupMemberships](#) – Lists the users in the specified Quick group.
- [ListUserGroups](#) – Lists the Quick groups that a Quick user is a member of.

Or you can ask your network administrator to query your identity provider to get this information.

The next two procedures provide instructions on how to finish creating a dataset for dynamic default values. The first procedure is for creating a dataset for a single-value DDP. The second one is for creating a dataset for a multivalue DDP.

To create a dataset for a single-value DDP

1. Create dataset columns with single-value parameters. The first column in the query or file should be for the people using the dashboard. This field can contain user names or group names. However, support for groups is only available in Quick Enterprise edition.
2. For each field that displays a dynamic default for a single-value parameter, add a column to the dataset. The name of the column doesn't matter—you can use the same name as the field or parameter.

Single-value parameters only work as specified if the combination of user entity and dynamic default is unique for that parameter's field. If there are multiple values a default field for a user entity, the single-value control for that field displays the static default instead. If no static default is defined, the control doesn't display a default value. Be careful if you use group names, because some user names can be members of multiple groups. If those groups have different default values, then this type of user name functions as a duplicate entry.

The following example shows a table that appears to contain two single-value parameters. We make this assumption because no user name is paired with multiple default values. To make this table easier to understand, we add the word 'default' in front of the field names from the analysis. Thus, you can read the table by making the following statement, changing the values for each row: When viewed by anacarolinasilva, the controls display a default region NorthEast and a default segment SMB.

Viewed-by	Default-region	Default-segment
anacarolinasilva	NorthEast	SMB
liujie	SouthEast	SMB
saanvisarkar	NorthCentral	SMB
zhangwei	SouthCentral	SMB

3. Import this data into Quick, and save it as a new dataset.
4. In your analysis, add the dataset that you created. The analysis needs to use at least one other dataset that matches the columns you defined for the defaults. For more information, see [Adding a dataset to an analysis](#).

To create a dataset for a multivalue DDP

1. Create dataset columns with multivalue parameters. The first column in the query or file should be for the people using the dashboard. This field can contain user names or group names. However, support for groups is only available in Quick Enterprise edition.
2. For each field that displays a dynamic default for a multivalue parameter, add a column to the dataset. The name of the column doesn't matter—you can use the same name as the field or parameter.

Unlike single-value parameters, multivalue parameters allow multiple values in the field that's associated with the parameter.

The following example shows a table that appears to contain a single-value parameter and a multivalue parameter. We can make this assumption because each user name has a unique value in one column, and some user names have multiple values in the other column. To make this table easier to understand, we add the word 'default' in front of the field names from the analysis. Thus, you can read the table by making the following statement, changing the values for each row: When viewed-by is liujie, the controls display a default-region value of SouthEast, and a default-city value of Atlanta. And if we read ahead one row, we see that liujie also has Raleigh in default-city.

Viewed-by	Default-region	Default-city
anacarolinasilva	NorthEast	New York
liujie	SouthEast	Atlanta
liujie	SouthEast	Raleigh
saanvisarkar	NorthCentral	Chicago
zhangwei	SouthCentral	Dallas
zhangwei	SouthCentral	Kansas City

In this example, the parameter that we apply `default-region` to works correctly whether it's a single-value or multivalue parameter. If it's a single-value parameter, two entries work for one user because both entries are the same value, `SouthEast`. If it's a multivalue parameter, it still works, except that only one value is selected by default. However, if we change the parameter that's using `default-city` as its default from a multivalue to a single-value parameter, we don't see these defaults selected. Instead, the parameter uses the static default, if there is one defined. For example, if the static default is set to `Atlanta`, `liujie` has `Atlanta` selected in that control, but not `Raleigh`.

In some cases, your static default value might also be used as a dynamic default. If so, make sure to test the control for a user name that doesn't use a default value that can be both.

If a user name belongs to multiple groups, the named user sees a set of default values that is a union of the two groups' default values.

3. Import this data into Quick, and save it as a new dataset.
4. In your analysis, add the dataset that you created. The analysis needs to use at least one other dataset that matches the columns you defined for the defaults. For more information, see [Adding a dataset to an analysis](#).

Use the following procedure to add a dynamic default parameter to your analysis. Before you begin, make sure that you have a dataset that contains the dynamic defaults for each user name

or group name. Also make sure that your analysis is using this dataset. For help with these requirements, see the procedures preceding.

To add a DDP to your analysis

1. In the Quick console, choose the **Parameters** icon at the top of the page and choose an existing parameter. Choose **Edit parameter** from the parameter's menu. To add a new parameter, choose the plus (+) sign near **Parameters**.
2. Choose **Set a dynamic default**.
3. Configure the following options with your settings:
 - **Dataset with default values and user information** – Choose the dataset that you created and added to your analysis.
 - **User name column** – To create defaults that are based on user names, choose the column in the dataset that contains the user names.
 - **Group name column** – To create defaults that are based on group names, choose the column in the dataset that contains the group names.
 - **Column for default value** – Choose the column that contains default values for this parameter.
4. Choose **Apply** to save your setting changes, and then choose **Update** to save the parameter changes. To exit without saving changes, choose **Cancel** instead.
5. Add a filter for each field that contains dynamic defaults to make the defaults work. To learn more about using filters with parameters, see [Using filters with parameters in Amazon Quick](#)

Amazon Quick uses the static default value for anyone whose user name doesn't exist in the dataset, doesn't have a default assigned, or doesn't have a unique default. Each person can have only one set of defaults. If you don't want to use dynamic defaults, you can set a static default instead.

Connecting to parameters in Amazon Quick

Use this section after you have a parameter set up, to connect it and make it work.

After you create a parameter, you can create consumers of the parameters. *Parameter consumers* are components that consume the value of a parameter, such as filters, controls, calculated fields, or custom actions.

You can navigate to each of these options in another way, as follows:

- To create a filter, choose the **Filter** icon at the top of the page. In short, you create a **Custom Filter** and enable **Use parameters**. The list shows only eligible parameters.
- To add a new control for the parameter, choose the **Parameters** icon at the top of the page. In short, choose your parameter, and then **Add control**.
- To use a parameter in a calculated field, either edit an existing calculated field, or add a new one by choosing **Add** at the top left. The parameter list appears below the field list.

 Note

You can't use multivalue parameters with calculated fields.

- To create a URL action, choose the **v**-shaped menu on a visual, and then choose **URL Actions**.

For more information on each of these topics, see the following sections.

Topics

- [Using filters with parameters in Amazon Quick](#)
- [Using calculated fields with parameters in Amazon Quick](#)
- [Using custom actions with parameters in Amazon Quick](#)
- [Using parameters in a URL](#)
- [Using parameters in titles and descriptions in Amazon Quick](#)

Using filters with parameters in Amazon Quick

Use this section to filter the data in an analysis or dashboard by a single-value parameter value. To use a multivalued parameter—one with a multiselect drop-down control—create a custom filter that is equal (or not equal) to the values.

Before using a filter with a parameter, you should already know how to work with filters.

1. Verify that your analysis has a parameter already created. Choose **Edit** from either the parameter or the control menu to find out what settings are in use.
2. Choose the **Filter** pane from the left of the screen. If there is already a filter for the field that you want to use, choose it to open its settings. Otherwise, create a filter for the field that you want to filter by parameter.

3. Choose **Use Parameters**.
4. Choose your parameters from the list or lists below **Use parameters**. For text (string) fields, first choose **Custom Filter**, and then enable **Use Parameters**.

For date fields, choose the **Start date** and **End date** parameters, as shown in the following screenshot.

For fields with other data types, choose **Select a parameter** and then choose your parameter from the list.

 Note

Parameters that can hold multiple values must use equal or not equal as the comparison type.

5. Choose **Apply** to save your changes.

Test your new filter by choosing the control near the top of the analysis. In this example, we use a basic parameter that has no defaults, and a dynamic control that is linked to the **Region** field in the sample dataset named **Sales Pipeline**. The control queries the data, returning all values.

If you delete or recreate a parameter that you are using in a filter, you can update the filter with the new parameter. To do this, open the filter, choose the new parameter that you want to use, and then choose **Apply**.

If you rename a parameter, you don't need to update the filter or any other consumers.

Using calculated fields with parameters in Amazon Quick

You can pass the value of a parameter to a calculated field in an analysis. When you create a calculation, you can choose existing parameters from the list of parameters under **Parameter list**. You can't create a calculated field that contains a multivalued parameter—those with a multiselect drop-down control.

For the formula, you can use any of the available functions. You can pass the viewer's selection from the parameter control, to the `ifElse` function. In return, you get a metric. The following shows an example.

```
ifelse(
```

```
`${KPIMetric}` = 'Sales',sum({Weighted Revenue})),  
  
`${KPIMetric}` = 'Forecast',sum({Forecasted Monthly Revenue})),  
  
`${KPIMetric}` = '# Active', distinct_count(ActiveItem),  
  
NULL  
  
)
```

The preceding example creates a metric (a decimal) that you can use in a field well. Then, when a user chooses a value from the parameter control, the visual updates to reflect their selection.

Using custom actions with parameters in Amazon Quick

A *custom action* enables you to launch URLs or filter visuals by selecting a data point in a visual or choosing the action name from the context menu. When you use a URL action with a parameter, you can pass or send parameters dynamically to the URL. To make this work, you set up a parameter, and then use it in the URL when you create a custom action with an action type of **URL action**. The parameters on both the sending and the receiving end must match in name and data type. All parameters are compatible with URL actions.

For details on creating a URL action, see [Creating and editing custom actions in Amazon Quick Sight](#). If you just want to use a parameter in a link without creating a URL action, see [Using parameters in a URL](#).

Using parameters in a URL

You can use a parameter name and value in a URL in Amazon Quick to set a default value for that parameter in a dashboard or analysis.

The following example shows the URL of a dashboard that sets a parameter for another dashboard.

```
https://us-east-2.quicksight.aws.amazon.com/sn/dashboards/abc123-abc1-abc2-abc3-  
abcdefef1234#p.myParameter=12345
```

In the previous example, the first part is the link to the target dashboard: `https://us-east-2.quicksight.aws.amazon.com/sn/dashboards/abc123-abc1-abc2-abc3-`

abcdefef1234. The hash sign (#) follows the first part to introduce the *fragments*, which contain the values that you want to set.

The values in the fragments aren't received or logged by AWS servers. This functionality keeps your data values more secure.

The fragment after # follows these rules:

- Parameters are prefixed with p.. The names are the parameter name, not the control name. You can view the parameter name by opening the analysis, and choosing **Parameter** on the left sidebar.
- The value is set using equals (=). The following rules apply:
 - Literal values don't use quotation marks.
 - Spaces inside values are automatically encoded by the browser, so you don't need to use escape characters when manually creating a URL.
 - To return all values, set the parameter equal to "[ALL]".
 - To assign the parameter's value to null, set it equal to %00. For example, p.population=%00.
 - In custom actions, target parameter names begin with \$, for example: <<\$passThroughParameter>>
 - In custom actions, parameter values display with angle brackets << >>, for example <<dashboardParameter1>>. The dashboard user sees the lookup value, not the variable.
- For a custom URL action, multivalue parameters only need one instance of the same parameter in the fragment, for example: p.city=<<\$city>>
- For a direct URL, multiple values for a single parameter have two instances of the same parameter in the fragment. For an example, see following.
- Ampersands (&) separate multiple parameters. For an example, see following.

The server converts the date to UTC and sends it to the backend as a string without a time zone. To use Universal Coordinated Time (UTC) dates, exclude the time zone. Following are some examples of date formats that work:

- 2017-05-29T00%3A00%3A00
- 2018-04-04 14:51 -08:00
- Wed Apr 04 2018 22:51 GMT+0000

```
https://us-east-2.quicksight.aws.amazon.com/sn/dashboards/abc123-  
abc1-abc2-abc3-abcdefef1234#p.shipdate=2018-09-30 08:01&p.city=New  
York&p.city=Seattle&p.teamMember=12&p.percentageRank=2.3
```

In the browser, this code becomes the following.

```
https://us-east-2.quicksight.aws.amazon.com/sn/dashboards/abc123-  
abc1-abc2-abc3-abcdefef1234#p.shipdate=2018-09-30%2008:01&p.city=New  
%20York&p.city=Seattle&p.teamMember=12&p.percentageRank=2.3
```

The previous example sets four parameters:

- `shipDate` is a date parameter: Sept 30, 2018.
- `city` is a multivalued string parameter: New York, and Seattle
- `teamMember` is an integer parameter: 12.
- `percentageRank` is a decimal parameter: 2.3.

The following example shows how to set values for a parameter that accepts multiple values.

```
https://us-east-2.quicksight.aws.amazon.com/sn/dashboards/abc123-abc1-abc2-abc3-  
abcdefef1234#p.MultiParam=WA&p.MultiParam=OR&p.MultiParam=CA
```

To pass values from one dashboard (or analysis) to another dashboard based on the user's data point selection, use custom URL actions. If you choose, you can also generate these URLs manually, and use them to share a specific view of the data.

For information on creating custom actions, see [Using custom actions for filtering and navigating](#).

Using parameters in titles and descriptions in Amazon Quick


When you create parameters in Amazon Quick, you can use them in titles and descriptions throughout your charts and analyses to dynamically display parameter values.

You can use parameters in the following areas of your analysis:

- Chart titles and subtitles

- Axis titles
- Legend titles
- Parameter control titles
- Sheet titles and descriptions

The following image shows a chart title that uses a parameter.



Use the following procedures to learn how to add parameters to areas throughout your analysis. For more information about parameters and how to create them, see [Parameters](#).

Adding parameters to chart titles and subtitles

Use the following procedure to learn how to add parameters to chart titles and subtitles.

To add a parameter to a chart title or subtitle

1. Open the **Properties** pane for the visual that you want to format.
2. In the **Properties** pane, choose the **Title** tab.
3. Select **Show title** or **Show subtitle**. These options might already be selected.
4. Choose the three dots at the right of **Edit title** or **Edit subtitle**, and then choose a parameter from the list.

The parameter is added to the title in the **Properties** pane. In the chart, the parameter value is displayed in the title.

For more information about editing titles and subtitles in visuals, see [Titles and subtitles on visual types in Quick](#).

Adding parameters to axis titles

Use the following procedure to learn how to add parameters to axis titles.

To add a parameter to an axis title

1. Open the **Properties** pane for the visual that you want to format.
2. In the **Properties** pane, choose the axis that you want to format.
3. Select **Show title**.
4. Choose the three dots at the right of the default axis title, and then choose a parameter from the list.

The parameter is added to the axis title in the **Properties** pane. In the chart, the parameter value is displayed in the axis title.

For more information about editing axis titles, see [Axes and grid lines](#).

Adding parameters to legend titles

Use the following procedure to learn how to add parameters to legend titles.

To add a parameter to a legend title

1. Open the **Properties** pane for the visual that you want to format.
2. In the **Properties** pane, choose **Legend**.
3. Select **Show legend title**.
4. Choose the three dots at the right of **Legend title**, and then choose a parameter from the list.

The parameter is added to the legend title in the **Properties** pane. In the chart, the parameter value is displayed in the legend title.

For more information about formatting legends, see [Legends on visual types in Quick](#).

Adding parameters to control titles

Use the following procedure to learn how to add parameters to parameter control titles.

To add a parameter to a parameter control title

1. Select the parameter control that you want to edit, choose the three dots at the right of the parameter control title, and then choose **Edit**.
2. In the **Edit control** page that opens, select **Show title**.
3. Choose the three dots at the right of **Display name**, and then choose a parameter from the list.

The parameter is added to the parameter control title.

For more information about using parameter controls, see [Parameter controls](#).

Adding parameters to sheet titles and descriptions

Use the following procedure to learn how to add parameters to sheet titles and descriptions in your analysis.

To add a parameter to a sheet title or description

1. On the analysis page, choose **Sheets** in the application bar and then choose **Add title** or **Add description**.

A sheet title or description appears on the sheet.

2. For **Sheet title** or for **Description**, choose the three dots at right, and then choose a parameter from the list.

The parameter is added to the sheet title or description and the parameter value appears in the text when you close the text box.

For more information about adding sheet titles and descriptions, see [Adding a title and description to an analysis](#).

Using custom actions for filtering and navigating

To add interactive options for dashboard subscribers (Quick readers), you create custom actions on one or more visuals in your analysis. Enhancing dashboards with custom actions helps people explore data by adding more context from within the dataset. It can make it easier to drill into the details and to find new insights in the same dashboard, a different dashboard, or a different application. You can add up to 10 custom actions to each visual in a dashboard.

Before you begin, it's helpful to do some planning. For example, identify fields that are good candidates for filtering, for opening a different sheet, for opening a URL, or for sending email. For each sheet, identify the widgets that display these fields. Then decide which widgets are going to contain actions. It's also a good idea to create a naming scheme so the names of the actions are consistent throughout the entire analysis. Consistent names make it easier for the person using your analysis to figure out what the action will do, plus they make it easier for you to maintain actions that you might be duplicating throughout the analysis.

Actions only exist on the dashboard widget where you create them and they work in the context of that widget's parent sheet and child fields that it displays. You can create actions only on specific types of widget: visuals and insights. You can't add them to other widgets, for example filter or list controls. Custom actions can only be activated from the widget where you create them.

To activate an action, the person using the analysis can left-click (select) or right-click (use the context menu) on a data point. A *data point* is an item in the dataset, for example a point on a line chart, a cell in a pivot table, a slice on a pie chart, and so on. If the person clicks a visual element, the *select* action is activated. This is the action that is currently a member of the **On select** category of the **Actions** in an analysis. If the person instead right-clicks a visual element, they can choose from a list of *menu* actions. Any action listed is currently a member of the **Menu option** category of the **Actions** in an analysis. The **On select** category can contain one and only one member action.

By default, the first action you create becomes the select action—the one activated by left-clicking. To remove an action from the **On select** category, change the action's **Activation** setting to **Menu option**. After you save that change, you can set a different action's **Activation** setting to **Select**.

You can choose from three **Action types** when you configure an action:

- **Filter action** – Filter data included in visual or in the entire sheet. By default, filters are available for all fields in the parent visual. Cascading filters are enabled by default. Filter actions work across multiple datasets by using automatically generated field mappings.

If the analysis uses more than one dataset, you can view the automatically generated field mappings for fields that exist in multiple datasets. To do this, choose **View field mapping** at the end of the action settings, while you're editing an action. If you are viewing a list of actions, choose **View field mapping** from the menu for each action. The field mappings appear in a new screen that shows the mapping between the initial dataset and all the other datasets in the visual. If no fields are automatically mapped, a message displays with a link to [Mapping and Joining Fields](#).

- **Navigation actions** – Enable navigation between different sheets in the same analysis.
- **URL actions** – Open a link to another web page. If you want to open a different dashboard, use a URL action. You can use a URL action to send data points and parameters to other URLs. You can include any available field or parameter.

If the URL uses the mailto scheme, running the action opens your default email editor.

Topics

- [Adding one-click interactive filters](#)
- [Creating and editing custom actions in Amazon Quick Sight](#)
- [Repairing custom actions](#)
- [Understanding field mapping for custom actions in Amazon Quick Sight](#)

Adding one-click interactive filters

One-click interactive filtering provides point-and-click filtering that cascades from the clickable visual to all the other visuals and insights on a sheet. Add this to your analysis to start with summaries and drill down into the metrics, all within the same dashboard sheet.

After you set this up, when you click a data point (for example, a point in a line chart), you instantly filter using all mapped fields on all the other visuals on that sheet. If you have multiple datasets, all target fields must be mapped for this to work. Also, you can only have one action that works by clicking a data point; all other actions work from the context menu.

Use the following procedure to create a one-click filter in an analysis.

To create a one-click filter on a visual or insight

1. In your analysis, choose a visual or insight that you want to add interactive filtering to.

2. Choose **Actions** from the Menu options dropdown in the upper right corner.
3. Choose **Filter same-sheet visuals**. Doing this immediately adds one-click filtering.
4. Repeat this process for each visual that you wish to make interactive.

Creating and editing custom actions in Amazon Quick Sight

You create one action for each task that you want to be able to add to a visual. The actions you create become part of the functionality of each visual or insight.

The following table defines when to use each type of action.

Action to perform	Type of action
Add or customize an interactive filter action, including one-click filters	Filter action
Open another sheet in the same dashboard	Navigation action
Open a sheet in a different dashboard in the same AWS account	URL action
Open a URL (https, http)	URL action
Send an email (mailto)	URL action

You can set the following attributes and options for a custom action:

- **Action name** – This is a descriptive name that you choose for the action. By default, actions are named **Action 1**, **Action 2**, and so on. If your custom action is activated from a context menu, this name displays in the menu when you right-click on a data point.

To make the action name dynamic, you can parameterize it. Use the plus icon near the action name header to display a list of available variables. Variables are enclosed in angle brackets << >>. Parameters are prefixed with a \$, for example <<\$parameterName>>. Field names have no prefix, for example <<fieldName>>.

- **Activation** – Available options are **Select** or **Menu option**. To use an action, you can *select* the data point (left-click) or navigate to the *menu option* in the context menu (right-click). Navigation

actions and URL actions listed in the middle of the context menu, just above **Color** options. Actions that are activated by menu are also available from the legend on a visual.

- **Action type** – The type of action that you want. Settings that are specific to an action type only display after you choose the action type.
- **Filter action** settings include the following:
 - **Filter scope** – The fields to filter on. To filter on all fields, choose **All fields**. Otherwise, choose **Select fields** and then turn off the items you don't want to target.

The default is **All fields**.

- **Target visuals** – The dashboard widgets to target. To apply the filter to all of them, choose **All visuals**. Otherwise, choose **Select visuals** and then turn off the items you don't want to target. When you apply a filter action to other visuals, the effect is called *cascading filters*.

The default is **All visuals**.

A cascading filter applies all the visuals that are set up in the **Target visuals** section of a specific filter action. Amazon Quick Sight initially evaluates your visuals and preconfigures the settings for you. But you can change the defaults if you wish to do so. You can set up multiple cascading filters on multiple visuals in the same sheet or analysis. When you are using the analysis or dashboard, you can use multiple cascading filters at the same time, although you activate each of these one at a time.

A filter action requires at least one target visual, because a filter action requires a source and a target. To filter only the current visual, create a regular filter instead by choosing **Filter** at left.

- **Navigation action** settings include the following:
 - **Target sheet** – The sheet to target.
 - **Parameters** – The parameters to send to the target sheet. Choose the plus icon to add an existing parameter.
- **URL action** settings include the following:
 - **URL** – The URL to open. URL actions can be deep links into another application. Valid URL schemes include https, http, and mailto.
 - **+ (Values)** – (Optional) The parameters to send to the target URL. Parameter names start with a \$. The parameters on both the sending and the receiving end must match in name and data type.

- **Open in** – Where to open the URL. You can choose **New browser tab**, **Same browser tab**, or **New browser window**.

Some types of actions enable you to include values from parameters or fields that are available in the visual or insight. You can type these in manually or choose + to select from a list. For the custom action to work, every field and parameter it references must be actively in use in the parent widget.

Use the following procedure to create, view, or edit a custom action in an analysis.

To create, view, or edit a custom action

1. With your analysis open, choose **Actions** from the **Menu options** dropdown in the upper right corner.

The existing actions, if any, display by activation type. To turn an existing action on or off, use the box to the right of the action's name.

2. (Optional) To edit or view an existing action, choose the menu icon next to the name of the action.

To edit the action, choose **Edit**.

To delete it, choose **Delete**.

3. To create a new action, choose either one of the following:
 - The add icon near the **Actions** heading
 - The **Define a custom action** button
4. For **Action name**, define an action name. To make the action name dynamic, use the plus icon to add parameter or field values.
5. For **Activation**, choose how the action runs.
6. For **Action type**, choose the action type you want to use.
7. For a **Filter action**, do the following:
 - a. For **Filter scope**, choose the scope of the filter.
 - b. For **Target visuals**, choose how far the filter cascades
8. For a **Navigation action**, do the following:

- a. For **Target sheet**, choose the target sheet.
 - b. For **Parameters**, choose the plus icon near the **Parameters** heading, select a parameter, and then choose a parameter value. You can choose all values, enter custom values, or select specific fields.
9. For a **URL action**, do the following:
- a. For **URL**, enter the hyperlink.
 - b. Choose the plus icon near the **URL** heading. Then, add variables from the list.
 - c. For **Open in**, choose how to open the URL.
10. After you are finished with the action, choose one of the following at the bottom of the **Actions** panel (you might need to scroll down):
- **Save** – Save your selections, and create the custom action.
 - **Close** – Close this custom action and discard your changes.
 - **Delete** – Delete this action.

Repairing custom actions

For a custom action to work, every field and parameter it references must be active in the parent widget. If a field is missing from the source widget, or if a parameter is missing from the analysis, the action for that field or parameter becomes unavailable. Menu actions are no longer included in the context menu. Select actions no longer respond to attempts to interact. However, in all other ways, the widget continues to function. No error displays to your users. You can fix broken filter actions and URL actions by adding the missing fields back to the broken visual or insight.

The following procedure explains how to fix an action that broke because someone removed a field or parameter without updating the action. These steps provide basic guidance how to fix this issue. However, use your own judgment on how or if you should make changes to the analysis. If you're not sure, it's better to ask a Amazon Quick administrator for assistance before you change anything. For example, there might be a way to restore a previous version of the analysis, which might be safer if you aren't sure what happened to it.

To remove a field from a broken action

1. From the start page, choose **Analyses**. Then choose the analysis to fix.

2. Choose the visual or insight where the action no longer works. Make sure that it's highlighted on the sheet.
3. Choose **Actions** from the Menu options dropdown in the upper right corner.
4. Locate the action you want to fix and choose **Edit**.
5. If the action type is **Filter action**, and you see an error that says *the field used by this action was removed*, check the settings for **Filter scope**. **Selected fields** can only display fields that are in the visual. To disable selected fields that are removed, choose one of the following:
 - Change the **Filter scope** setting to **All fields**. Doing this enables the widget to filter on every field.
 - If you want to use a list of **Selected fields**, verify the list of fields. If you need to include another field, you need to add it to the visual first.
6. If the action type is **Navigation action**, follow the guidance on the error message, which reflects the type of change that caused the error.
7. If the action type is **URL action**, check the **URL** setting for variables marked with double angle brackets (<<FIELD-OR-\$PARAMETER>). Open the list of available variables by choosing the plus icon. Remove any fields or parameters that aren't in the list. Be sure you also remove the matching *URL parameter* and its separator (? for the first URL parameter, or & for subsequent parameters). The following examples show (in **bold**) which part is removed if you were removing the field named **Product** from the visual.

```
https://www.example.com/examplefunction?q=<<Product>
```

```
https://www.example.com/examplefunction?q=<<Product>&uact=<<$CSN>
```

```
https://www.example.com/examplefunction?pass=yes&q=<<Product>+<<City>&oq=<<Product>+<<City>&uact=<<$CSN>
```

Make sure to test the new URL.

8. (Optional) To delete the action, scroll to the end and choose **Delete**.
9. When you are finished, confirm your changes to the action. Scroll to the bottom of the **Action** pane and choose **Save**.

If the error also exists in an associated dashboard, share and publish the dashboard again to propagate the fix.

Understanding field mapping for custom actions in Amazon Quick Sight

Automated field mapping is based on identical fields. Fields with the same name and data type map automatically across datasets. Their field names and data types must be an exact match. This works similar to a join, except that it is automatically generated based on names and data types for every matching field. If you are missing fields, you can create them by using calculated fields in the dataset that's missing a field. If you don't want to have some of the fields mapped to each other, you can rename or remove them from the dataset.

It's important to make sure that all target fields are mapped if they are enabled for use with a filter action (in the **Filter scope**). Doing this allows filtering to apply automatically. If some target fields aren't mapped, the automatic filtering doesn't work.

Mapping is generated only when you create or save a custom action. So after every change that affects the mapping, make sure to return to it and save it again. When you create an action, mapping is based on the fields as they exist at that point. When you save an action, any mapped fields that you renamed since you created the custom action stay mapped. However, if you alter the data type of a mapped field, the mapping is removed.

If your mapping is missing some fields, you can do one of the following to fix it:

- Only target the mapped fields, by removing the unmapped fields from the **Filter scope**.
- Remove the visual in question from the target visuals.
- Create calculated fields to supply the missing fields for the mapping, and then save your custom action.
- Edit the dataset and rename the fields or change their data types, and then save your custom action.
- Edit the dataset and rename the fields or change their data types, and then resave your custom action.

Note

The information that displays on the mapping screen shows the configuration from the most recent time you saved it. To refresh or update the view, save the action again.

If you add or edit datasets, they aren't automatically mapped or remapped. This causes the filtering to work incorrectly. For example, suppose that you add a new dataset, then create visuals for it.

The new visuals won't respond to filter actions, because there is no field mapping to connect them. When you make changes, remember to save your custom actions again to redo the field mappings.

If you remove a parameterized field or any other targeted field from the source visual, the action that uses it breaks. The action for the missing field either doesn't work when you select a data point, or it's hidden from the context menu.

For information about preparing your dataset for automated field mapping, see [Mapping fields](#).

Working with pixel perfect reports in Amazon Quick Sight

With Amazon Quick Sight Pixel Perfect Reports, you can create, schedule, and share highly formatted multipage PDF reports. You can also schedule data exports as CSV files using Quick Sight's existing web interface. This unifies historically separate systems for dashboards and reports.

Report creators can use Quick Sight's browser-based authoring experience to connect to a broad range of supported data sources and create highly formatted reports. They can specify the exact page size, length, and arrangement of images, charts, and tables with pixel-level precision. Authors can then use Quick Sight's scheduling mechanisms to set up and schedule highly personalized report delivery to end users, or archive reports for future use.

Pixel perfect reports are designed to be printed or distributed. Pixel perfect report content is formatted to fit paper sizes and it displays all the data in a table and pivot table, even if the data spans multiple pages. They are formatted for exact paper sizes and you can control page layout exactly. Each pixel perfect report can generate a PDF of up to 1,000 pages.

Pixel perfect reports provide all available data that is present when the report is published as a PDF or CSV. For example, let's say you have a table with 10,000 rows. A pixel perfect report presents the entire report across multiple pages for readers to view in its entirety. If you include this same table in an interactive dashboard report, the generated PDF includes a snapshot of the table that fills in a single page that can be scrolled through. These customized reports can be sent out in email bursts that generate up to thousands of personalized PDF or CSV reports to individual users and groups.

Note

Pixel perfect reports are not available in the eu-central-2 Europe (Zurich) region.

Topics

- [Getting started](#)
- [Creating reports from an analysis in Amazon Quick Sight](#)
- [Formatting reports in Amazon Quick Sight](#)
- [Consuming pixel perfect reports in Amazon Quick Sight](#)
- [Unsubscribe from paginated reporting in Quick Sight](#)

Getting started

To get started working with Amazon Quick Sight pixel perfect reports, first get the pixel perfect reporting add-on for your Amazon Quick account. Pricing for the add-on applies to your entire Quick account and isn't specific to a Region. After you subscribe to Quick reporting, authors can begin creating, scheduling, and sending pixel perfect reports.

Get the Quick pixel perfect reports add-on

Before you can work with pixel perfect reports in Amazon Quick Sight, you must add the **Pixel-Perfect Reports add-on** to your Quick subscription.

To get the pixel perfect reporting add-on in Amazon Quick Sight

1. On the Quick start page, choose your user name at the upper right, and then choose **Manage Quick**.
2. Choose **Manage subscriptions**, and then choose **Pixel-Perfect Reports**.
3. Choose the subscription plan that you want. You can choose between a monthly plan and an annual plan.
4. Review the Pixel-Perfect Reports add-on pricing information and then choose **Confirm subscription**.

After you get the pixel perfect reports add-on, it may take several minutes for your subscription to take effect. When your subscription takes effect, you will be able to begin creating pixel perfect reports in Amazon Quick Sight.

Creating reports from an analysis in Amazon Quick Sight

Pixel perfect reports are created at the sheet level of an analysis in Amazon Quick Sight. When you create a new analysis or a new sheet in an existing analysis, you choose whether to make the

new sheet an **Interactive dashboard** or a **Pixel perfect report**. This way, you can have analyses for interactive dashboards only, analyses for pixel perfect reports only, or you can have an analysis that includes both interactive dashboards and pixel perfect reports.

There are three ways to create a pixel perfect report. You can create a new report from a new sheet in an analysis, you can duplicate an interactive sheet in a dashboard, or you can duplicate a pixel perfect report that already exists. Use the procedures below to create a pixel perfect report.

Creating reports from an analysis in Amazon Quick Sight

To create a pixel perfect report from a new analysis

1. On the Quick homepage, choose **Analyses**, and then choose **New analysis**.
2. Choose the dataset that you want to include in your new analysis, and then choose **USE IN ANALYSIS** in the top right.
3. In the **New sheet** pop-up that appears, choose **Pixel perfect report**.
4. (Optional) Choose the paper size that you want for your pixel perfect report. You can choose from the following options:
 - US letter (8.5 x 11 in.)
 - US legal (8.5 x 14 in.)
 - A0 (841 x 1189 mm)
 - A1 (594 x 841 mm)
 - A2 (420 x 594 mm)
 - A3 (297 x 420 mm)
 - A4 (210 x 297 mm)
 - A5 (148 x 210 mm)
 - Japan B4 (257 x 364 mm)
 - Japan B5 (182 x 257 mm)

The default paper size is US letter (8.5 x 11 in.)

5. (Optional) Choose between a portrait and landscape arrangement for the sheet. The default option is portrait.
6. Choose **ADD**.

If you want to create a new pixel perfect report in an existing analysis, choose the plus sign (+) icon to the right of the sheet tabs in your analysis and follow steps 3-6 from the preceding procedure.

Creating reports from an existing dashboard in Amazon Quick Sight

You can also create a pixel perfect report by duplicating an interactive sheet and converting the duplicate sheet into a pixel perfect report.

To create a pixel perfect report from an interactive sheet

1. From the sheet that you want to duplicate in an analysis, choose the dropdown next to the name of the sheet that you want to convert.
2. Choose **Duplicate to report**.

You can convert an interactive sheet to a pixel perfect report, but you can't convert a pixel perfect report to an interactive sheet.

Duplicate an existing report in Amazon Quick Sight

This section will go over how to copy a report.

To copy a pixel perfect report

1. From the sheet that you want to duplicate in an analysis, choose the dropdown next to the name of the sheet that you want to convert.
2. Choose **Duplicate**.

Formatting reports in Amazon Quick Sight

Use this section to learn how to format a pixel perfect report in Amazon Quick Sight.

Topics

- [Working with sections](#)
- [Changing paper size, margins, and orientation](#)
- [Adding and removing page breaks to a report](#)
- [Adding and deleting visuals to a report](#)
- [Adding a text box to a report](#)

- [Setting up prompts for paginated reports](#)

Working with sections

A *section* is a container for different visuals that grow vertically to contain contents. Each section is rendered to completion, one after the other, accommodating configured page breaks and section settings. Headers and footers are special types of section that have a predefined size, location, and repetition throughout each page of a report.

Each section in a pixel perfect report can be formatted independently from other sections in the report. Visuals can be dragged and dropped anywhere you want, similar to a free-form layout in an interactive sheet. Visuals can also be overlapped, resized, or brought forward or to the back of the section. Additionally, you can change the margins within a section to make the grouping of visuals stand out from the rest of the report.

Every report in Quick Sight needs at least one section. You can add multiple sections to group different sets of visuals together, or to control the rendering order for different groupings of visuals.

Each pixel perfect report sheet supports up to 30 sections, including headers and footers.

Use the topics listed below to learn more about sections.

Topics

- [Adding, moving, and deleting sections](#)
- [Headers and footers](#)
- [Section padding](#)
- [Create repeating sections](#)

Adding, moving, and deleting sections

Add a new section

To add a new section to a pixel perfect report, use the following procedure.

To add a new section to a pixel perfect report

1. From the Quick homepage, choose **Analyses** and then choose the analysis that contains the report that you want to add a section to.

2. Choose the sheet that contains the pixel perfect report that you want to add a section to.
3. Choose the **ADD (+)** icon in the top left corner, and choose **Add section**.

You can also add a section by choosing the plus (+) icon at the bottom of an existing section and choosing **Add section**.

When you choose **Add section**, a new section is added to the bottom of the report.

You can't create a section inside of another section. If you select an existing section and then choose **Add section**, a new section will appear at the bottom of the report.

When you have multiple sections in a pixel perfect report, they can be arranged in any order that you want.

Move a section

To move a section in a report

1. Choose the section that you want to move, and then choose the three-dot icon in the right corner to open the on-section menu.
2. Choose where you want to move your section. You can choose from the following options:
 - **Move section to top**
 - **Move section up**
 - **Move section down**
 - **Move section to bottom**

In some cases, you aren't able to select some of the preceding options. For example, if your section is already at the bottom of the report, you can't select **Move down** or **Move section to bottom**.

Sections are named according to their ascending order in the report. When you move a section up or down in a report, every section affected by the move is renamed according to the new ascending order.

When you delete a section from a pixel perfect report, the names of the remaining sections may change depending on where the deleted section was located. For example, say you decide to delete Section 1. When you delete the section, the previous Section 2 will move up the report and become the new Section 1.

Delete a section

To delete a section from a report

1. Navigate to the section that you want to delete and choose the three-dot icon in the upper right corner to open the on-section menu.
2. Choose **Delete**.

Headers and footers

Headers and *Footers* are optional special sections located at the top and bottom of a pixel perfect report. Headers and footers are commonly used to display basic information like the date the report was created or the page number. You can interact with headers and footers the same way you interact with a regular section in a report.

By default, every report in Amazon Quick Sight has a header and a footer. To remove the header or footer from your report, use the following procedure.

To remove a header or footer from a pixel perfect report

1. In your pixel perfect report, navigate to the header or footer that you want to delete and open the **On-section**.
2. Choose **Delete**.

When you delete a header or footer from your report, you are deleting the header or footer from every page of the report. You can't have a header or footer on some pages but not others.

If you have removed the header or footer from your report but want them to be visible again, use the following procedure.

To add a header or footer to a pixel perfect report

1. Navigate to the pixel perfect report that you want to add a header or footer to and choose **Insert** from the top menu.
2. Choose **Add header** or **Add footer**.

Section padding

You can use section padding to change the margins of an individual section in a pixel perfect report. By default, all sections in a report use the page margins that are configured and applied to the entire report. You can also add section padding to a header or footer. With section padding, you can make a section stand out from other sections by creating another set of margins. Apply the new set of margins to the section on top of the page margins that the rest of the report uses.

To change the section padding of a section

1. Navigate to the section that you want to add section padding to and open the **Edit section**.
2. In the **Padding** section of the **Edit section**, enter the amount of padding you want in inches. You can customize the padding of every side of the section (top, bottom, left, and right).

You can't use section padding to decrease the margins of the section. For example, if the margins of the entire pixel perfect report are 1 inch, you can only add to that value with section padding.

Create repeating sections

Use repeating sections to create duplicates of specific sections of a report to show one or more-dimension values. The data in the repeating section is sliced to match the dimensions of the section. Repeating sections can be replicated at scale to decrease the amount of time it takes to build reports.

Use the following procedures to create and configure a repeating section in a report.

To define a repeating section

1. Navigate to the section that you want to add a repeating behavior to and choose the **Edit repeating section** (triple panel).
2. In the **Edit section** pane that opens, choose **ADD DIMENSION**, and then choose the dimension that you want to add.
3. To add additional dimensions, repeat Step 2. You can add up to 3 dimensions in each repeating section configuration.

Considerations for repeating sections

The following limits apply to repeating sections.

- Insight visuals aren't supported for repeating sections.
- Repeating sections dimensions that are only from the last dataset that was selected for use in the analysis.

After you create a repeating section, you can define sorting and limits to the repeating section configuration. You can also use text boxes to add system parameters to repeating sections.

Define sorting in a repeating section

To define sorting in a repeating section

1. Navigate to the section that you want to add a repeating behavior to and choose the **Edit repeating section** (triple panel).
2. In the **Edit section** that opens, choose the ellipsis (three dots) next to the dimension that you want to change.
3. Navigate to the **Repeating** tab and choose the ellipsis (three dots) next to the dimension that you want to sort, and then choose **Edit**.
4. For **Sort by**, use the dropdown to choose the dimension that you want to sort by.
5. For **Aggregation** dropdown, choose the aggregation that you want to specify.
6. For **Sort order**, choose **Ascending** or **Descending**.

Define limits in a repeating section

You can set limits to show only a certain number of distinct dimension values for each dimension of a repeating section. You can choose to show between 1 and 1000 distinct values. The default limit is 50.

To define limits in a repeating section

1. Navigate to the section that you want to add a repeating behavior to and choose the **Edit repeating section** (triple panel).
2. In the **Edit section** that opens, choose the ellipsis (three dots) next to the dimension that you want to change.
3. For **Limit to**, enter the number of values that you want to limit the sorting to. You can enter a number between 1 and 1000.

Considerations for limits

The following limitations apply to limits in repeating sections.

- An *instance* is defined as a distinct value of a dimension or a unique combination of values of multiple dimensions.
- If the number of unique instances for a dimension in a repeating section exceeds 1000, the PDF report is NOT generated. If this occurs, try one of the following options.
 - Define a limit for your dimension.
 - Create a sheet level filter to restrict the dimension values.
 - Use row level security (RLS) to restrict the dimension values.
 - Apply dataset filters.

Add system parameters to repeating sections

You can use text boxes to add system parameters to your paginated report's repeating section. This makes it possible to access dimensions that have been used to configure repeating sections. Repeating sections and dimensions need to be configured before you can access the dimensions in a text box. The system parameters can only be used within a repeating section.

To add system parameters to a repeating section from a text box

1. Choose the text box visual that you want, and then choose the **System parameters** icon in the far right of the text box toolbar.
2. From the dropdown that appears, choose the parameter that you want.

Add page breaks to repeating sections

Similar to section page breaks, you can add page breaks to repeating sections.

To add a page break to a repeating section

1. Navigate to the section that contains the repeating behavior that you want to change and choose the **Edit repeating section** (triple panel) icon.
2. In the **Repeating** tab of the **Edit section** pane that appears, check the box titled **Page break after each instance**.

An instance is defined as a distinct value of a dimension or a unique combination of values of multidimensions. If you clear the **Page break after each instance** checkbox, the page break is removed.

Changing paper size, margins, and orientation

After you create a pixel perfect report in Amazon Quick Sight, you can change the report format, orientation, and margins from the **Analysis settings** menu whenever you want.

To change the paper size of a pixel perfect report

1. From the Quick homepage, choose **Analyses**, and then choose the analysis that contains the pixel perfect report that you want to change.
2. Choose **Sheets** in the file menu and select **Layout Settings**.
3. Open the **Paper size** dropdown menu and choose the paper size that you want. Choose from the following options:
 - US letter (8.5 x 11 in.)
 - US legal (8.5 x 14 in.)
 - A0 (841 x 1189 mm)
 - A1 (594 x 841 mm)
 - A2 (420 x 594 mm)
 - A3 (297 x 420 mm)
 - A4 (210 x 297 mm)
 - A5 (148 x 210 mm)
 - Japan B4 (257 x 364 mm)
 - Japan B5 (182 x 257 mm)
4. Choose **Apply**.

To change the orientation of a report

1. From the Quick homepage, choose **Analyses**, and then choose the analysis that contains the pixel perfect report that you want to change.
2. Choose the **Settings** icon on the left.
3. Choose the orientation for your report, and then choose **Apply**.

To change the margins of a report

1. From the Quick start page, choose **Analyses**, and then choose the analysis that contains the pixel perfect report that you want to change.
2. Choose **Edit < Analysis Settings**.
3. Enter the margin values that you want your report to have, and then choose **Apply**.

Margin values are applied to every page of a pixel perfect report. You can't set custom settings for specific pages in a report, but you can set custom margins for sections using section padding. For more information on section padding, see [Section padding](#). Margin values are expressed in inches. The default margins for all reports are 0.5 inches.

Adding and removing page breaks to a report

You can add page breaks between sections of a pixel perfect report to organize the way data is rendered when the report is published by page. For example, let's say you have a report that contains two sections that are each 2.5 pages long. By default, Section 2 begins on the third page of the report directly following the end of Section 1. If you add a page break to the end of the Section 1, Section 2 begins on a new page, even if the last page of Section 1 only uses half of a page. This is useful when you don't want different sections to share pages, but you don't know how many pages each section will need.

To add or delete a page break

1. Select your section and choose the **Edit section** icon in the top left corner.
2. In the **Edit section** pane that opens on the left, select the **Page break after** check box.
3. Choose **Apply**.

When you check the **Page break after** box, a page break will appear at the end of the section. If you remove the check from the **Page break after** box, the page break is removed from the end of the section. Also, the proceeding section renders directly under the last page of the section, even if it causes the two sections to share a page.

You can also add or remove a page break from a report by choosing the plus (+) icon at the bottom of an existing section and choosing **Add page break** or **Remove page break**.

Adding and deleting visuals to a report

To add visuals to a section in a pixel perfect report

1. In your pixel perfect report, select the section that you want to add a visual to.
2. Choose the **ADD (+)** icon in **Visuals** pane.
3. Choose the visual type that you want to use in your report.

After you add a visual to a report, you can interact with it the same way you would if the visual was part of an interactive dashboard. You can drag and drop visuals anywhere you want, similar to a free-form layout in a Quick Sight interactive dashboard sheet. You can also overlap visuals, resize them, or bring them forward or to the back of the section. For more information on formatting visuals in Amazon Quick Sight, see [Formatting in Amazon Quick](#).

To delete a visual

1. In the section that you want to delete a visual from, select the visual that you want to delete.
2. Choose the three-dot icon in the upper right corner of the visual to open the on-visual menu.
3. Choose **Delete**.

When you delete a visual from a section of a pixel perfect report, you are only deleting that specific visual from the report. Any duplicate visuals that are located in different sections of the report will remain in the report.

Adding a text box to a report

You can add text boxes to your pixel perfect reports to add context to your reports. Text box visuals can also be used boxes to add hyperlinks to external websites. To customize the font, font style, text color, text spacing, text alignment, and text size, use the text box toolbar that appears when you select the visual.

To add a text box to a report

1. In your pixel perfect report, select the section that you want to add a text box to.
2. Choose the **Text box** icon in the task bar.
3. The new text box appears in the section of the report that you selected.

To edit a text box, select the text box and begin typing what you want. A toolbar appears that you can use to make changes to the formatting and style of the text.

To delete a text box

1. In the section that you want to delete a text box from, select the text box that you want to delete.
2. Choose the three-dot icon in the upper right corner of the visual to open the on-text box menu.
3. Choose **Delete**.

Text box system parameters

Use text boxes to add system parameters to your pixel perfect report's headers and footers. Text box system parameters appear on the far right side of the text box toolbar. You can add the following parameters to a header or footer of your report:

- Page numbers: The current page number of the report.
- Report print date: The date the report was generated.

To add a page number parameter to your text box, choose the number (#) icon on the far right side of the text box toolbar. To add a `PrintDate` parameter to your text box, choose the calendar icon on the far right side of the text box toolbar.

For more advanced parameter options, add an insight to your paginated report.

Setting up prompts for paginated reports

Amazon Quick authors can create prompts on pixel-perfect reports to allow dashboard users to filter data in on-demand and scheduled reports. *Prompts* behave the same way a filter or control behaves in an interactive sheet.

To define a prompt in a pixel perfect report

1. On a pixel perfect sheet, define a filter control or a parameter control. For more information on filter controls to sheets, see [Adding filter controls to analysis sheets](#). For more information on parameter controls, see [Using a control with a parameter in Amazon Quick](#).
2. In the new filter or parameter, choose the prompt values that you want. The new prompts are immediately reflected on the sheet.

3. To export the report with the new prompts, choose **File**, and then choose **Export to PDF**.

Prompts can't be moved to the sheet itself. Instead, they are displayed on the top panel.

After a prompt is created for a pixel-perfect report and is published as a dashboard, Quick authors can use the new prompt to configure and schedule reports that are sent to Quick dashboard viewers. Dashboard viewers can also use these prompts to create their own scheduled reports. For more information about reader generated reports, see [Creating a reader generated report in Amazon Quick Sight](#).

Consuming pixel perfect reports in Amazon Quick Sight

When an Amazon Quick author publishes and then sends a scheduled pixel perfect report, Quick will generate and save a snapshot of the report that is sent out. Whenever you go to view the pixel perfect report's dashboard, you will see the generated snapshot of the most recently sent report. If you try to view your report's dashboard but you haven't sent an email report yet, you are prompted to schedule your first report to see the dashboard snapshot. For more information on scheduling an email report, see [Scheduling and sending Quick Sight reports by email](#).

If a Quick author has set up a prompted report for a Quick Sight pixel perfect report, Quick readers can use the prompt to schedule their own on-demand reports for themselves. For more information about reader-generated reports, see [Creating a reader generated report in Amazon Quick Sight](#). For more information about prompts for pixel perfect reports, see [Setting up prompts for paginated reports](#).

Users can't interact with a published pixel perfect report the same way they can interact with a published interactive sheet. Unlike interactive sheets, pixel perfect reports generate static snapshots of data that is presented in groups of visuals or text boxes. These static snapshots are generated at the time that the report is sent, so that the audience can see the latest version of the data in the report. Pixel perfect reports are especially useful for generating invoices or weekly business reviews. Users can then compare the current pixel perfect reports with reports that were generated in the past to better track their business data.

Viewing a report's snapshot history

Every time you send out a scheduled pixel perfect report, Amazon Quick saves a copy of the generated snapshot that is sent out for your reference. You can view these snapshots at any time in the Quick console.

To view a report's snapshot history

1. From the Quick homepage, choose **Dashboards**, and then choose the dashboard whose snapshot history you want to see.
2. Choose the **Scheduling** icon in the top right toolbar, and then choose **Recent snapshots**.
3. In the **Recent snapshots** pane that appears on the right, choose the snapshot to view, and then choose the download button next to the file that you want to download.

Unsubscribe from paginated reporting in Quick Sight

You can unsubscribe from Quick Sight pixel perfect reporting at any time. Once you unsubscribe from pixel perfect reporting, you will lose the ability to create and schedule pixel perfect reports in Quick Sight. You are still able to access your existing paginated reports, but you won't be able to make changes or schedule new reports.

To unsubscribe from pixel perfect reporting in Amazon Quick Sight

1. From any page in Quick, choose your user name in the upper right, and choose **Manage Quick**.
2. Choose **Manage subscriptions** on the left.
3. On the **Manage subscriptions** page, locate the **Pixel-Perfect Reports** section and choose **Manage**.
4. Scroll down to your chosen subscription plan and choose **Cancel subscription**.

Working with items on sheets in Amazon Quick Sight analyses

Use this section to learn how to work with visuals and other items as you author sheets in Quick Sight

Topics

- [Adding visuals to Quick Sight analyses](#)
- [Using Topics on sheets in Amazon Quick Sight](#)
- [Visual types in Amazon Quick Sight](#)
- [Formatting in Amazon Quick](#)
- [Customizing data presentation](#)

Adding visuals to Quick Sight analyses

A *visual* is a graphical representation of your data. You can create a wide variety of visuals in an analysis, using different datasets and visual types.

After you have created a visual, you can modify it in a range of ways to customize it to your needs. Possible customizations include changing what fields map to visual elements, changing the visual type, sorting visual data, or applying a filter.

Quick Sight supports up to 50 datasets in a single analysis, and up to 50 visuals in a single sheet, and a limit of 20 sheets per analysis.

You can create a visual in several ways. You can select the fields that you want and use AutoGraph to let Amazon Quick Sight determine the most appropriate visual type. Or you can choose a specific visual type and choose fields to populate it. If you aren't sure what questions your data can answer for you, you can choose **Suggested** on the tool bar and choose a visual that Amazon Quick Sight suggests. Suggested visuals are ones that we think are of interest, based on a preliminary examination of your data. For more information about AutoGraph, see [Using AutoGraph](#).

You can add more visuals to the workspace by choosing **Add**, then **Add visual**. Visuals created after June 21, 2018, are smaller in size, fitting two on each row. You can resize the visuals and drag them to rearrange them.

To create a useful visual, it helps to know what question you are trying to answer as specifically as possible. It also helps to use the smallest dataset that can answer that question. Doing so helps you create simpler visuals that are easier to analyze.

Fields as dimensions and measures

In the **Visuals** pane, dimension fields have blue icons and measure fields have orange icons. *Dimensions* are text or date fields that can be items, like products. Or they can be attributes that are related to measures and can be used to partition them, like sales date for sales figures. *Measures* are numeric values that you use for measurement, comparison, and aggregation. You typically use a combination of dimension and measure fields to produce a visual, for example sales totals (a measure) by sales date (a dimension). For more information about the types of fields expected by the different visual types, see the specific visual type topics in the [Visual types in Amazon Quick Sight](#) section. For more information about changing a field's measure or dimension setting, see [Setting fields as a dimensions or measures](#).

Field limitations

You can only use one date field per visual. This limitation applies to all visual types.

You can't use the same field for more than one dimension field well or drop target on a visual. For more information about how expected field type is indicated by field wells and drop targets, see [Using visual field controls](#).

Searching for fields

If you have a long field list in the **Fields list** pane, you can search to locate a specific field. To do so, choose the search icon at the top of the **Data** pane and then enter a search term into the search box. Any field whose name contains the search term is shown. Search is case-insensitive and wildcards aren't supported. Choose the cancel icon (X) to the right of the search box to return to viewing all fields.

Adding a visual

Use the following procedure to create a new visual.

To create a new visual

1. Open the [Quick console](#).
2. On the Quick homepage, choose the analysis that you want to add a visual to.
3. On the analysis page, choose the dataset that you want to use from the dataset list at the top of the **Data** pane. For more information, see [Adding a dataset to an analysis](#).
4. Open the **Visualize** pane, choose **Add**, and then choose **Add visual**.

A new, blank visual is created and receives focus.

5. Use one of the following options:
 - Choose the fields to use from the **Data** pane at left. If the fields aren't visible, choose **Visualize** to display it. Amazon Quick Sight creates the visual, using the visual type it determines is most compatible with the data you selected.
 - Choose the dropdown arrow next to the **ADD** button to choose a visual type. After the visual is created, choose the fields that you want to populate it.
 1. Choose the icon of a visual type from the **Visual types** pane.



The field wells display the fields that are visualized.



2. From the **Data** pane, drag the fields that you want to use to the appropriate field wells. Typically, you want to use dimension or measure fields as indicated by the color of the

target field well. If you choose to use a dimension field to populate a **Value** field well, the **Count** aggregate function is automatically applied to it to create a numeric value.

Amazon Quick Sight creates the visual using the visual type you selected.

- Create a visual using a suggestion.

On the tool bar, choose **Suggested**, then choose a suggested visual.

Importing Amazon Quick Sight visuals to an analysis

Quick Sight authors can import Quick Sight visuals from one analysis or dashboard to a new analysis that has access privileges. When you import a visual from a Quick Sight analysis or dashboard to another Quick Sight analysis, the following dependencies are imported along with the visual.

- Datasets associated with the visual
- All parameters that are configured to the visual
- Calculated fields that are configured to the visual
- Filter definitions
- Visual properties
- Conditional formatting rules

Use the following sections to learn more about importing Quick Sight visuals.

Topics

- [Considerations](#)
- [Import a visual](#)

Considerations

Before you import a visual, review the following limitations.

- The Quick Sight author that wants to import a visual must have ownership privileges to the analysis that they want to import the visual to
- Filter controls can't be imported
- Importing visuals from multiple sheets at a time is not supported

- Some user configurations including filter configurations that are maintained against bookmarks and alerts are not supported

Import a visual

Use the following procedure to import a visual from a source dashboard or analysis to a different analysis.

1. Open the [Quick console](#).
2. Open the analysis that you want to import a visual to.
3. Choose **File**, and then choose **Import**. Alternatively, you can choose the **Import** icon in the **ADD** toolbar.
4. The **Asset explorer** modal opens. A list of all eligible source analyses and dashboards that you can access are displayed. Choose the artifact that you want to import a visual from, and then choose **LOAD**. Alternatively, enter the name of the source artifact that contains the visual that you want to import in the **Find source to insert** search bar. Choose the artifact that you want, and then choose **LOAD**.
5. In the **Select visuals to import** page that opens, choose the sheet that contains the visuals that you want to import, and then choose the visuals that you want to import. You can only import visuals from one sheet at a time. When you have chosen all visuals that you want to import, choose **IMPORT**.

After a successful import job, the imported visuals are added to the destination analysis. The imported visuals retain the original properties that were configured to them in the source dashboard or analysis. Imported visuals inherit the theme-level properties from the theme that is applied to the destination analysis.

Duplicating Quick Sight visuals

You can duplicate a visual to make a new copy of it on the same sheet or on a different sheet.

To duplicate a visual, on the v-shaped on-visual menu, choose **Duplicate visual to**, then choose the sheet where you want the visual to appear. The display automatically shows you the duplicated visual.

Duplicated visuals keep all the same filters and settings as the source visual. However, if you duplicate a visual onto a different sheet, all of its copied filters apply to the duplicate only. All

copied filters are scoped down to apply only to that visual. If you want the filters to apply to more visuals on the new sheet, edit the filter and change the setting.

Parameters and controls apply to all sheets. To make parameter controls work with a visual that you duplicate to a different sheet, add filters on the target sheet and connect them to the parameter. To do this, choose **Custom filter** as the filter type.

Renaming Amazon Quick Sight visuals

Use the following procedure to rename a visual.

To rename a visual

1. On the analysis page, choose the visual that you want to rename.
2. Select the visual name at the top left of the visual and enter a new name.
3. Press **Enter** or click outside of the visual name field to save the new name.

Viewing visual data in Amazon Quick Sight

Amazon Quick Sight offers a variety of ways to see the details of the data being displayed in a visual. The axes or rows and columns of the visual (depending on the visual type) have labels. Hovering over any graphical element in a visual displays the data associated with that element. Some visual types use visual cues to emphasize the element that you are hovering over and make it easier to differentiate. For example, the visual type might change the color of the element or highlight it.


Use the following sections to learn more about viewing data in visuals.

Topics


- [Viewing visual details](#)
- [Scrolling through visual data](#)
- [Focusing on visual elements](#)
- [Excluding visual elements](#)
- [Searching for specific values in your data in Quick Sight](#)

Viewing visual details

When viewing a visual, you can hover your cursor over any graphical element to get details on that element. For example, when you hover over a single bar on a bar chart, information about that specific bar displays in a tooltip.



Hovering your cursor over a single data point on a scatter plot also displays information about that specific data point.



You can customize the information that appears when you hover your cursor over data in a chart. For more information, see [Tooltips](#).

Scrolling through visual data

For bar charts, line charts, and pivot tables, the content of the visual can be larger than the size that you want the visual to be.

In these cases, scrub bars appear so you can either reduce the data that is displayed or scrub through it. This process is similar to the way that you can scrub through a video.

To reduce the length of the scrub bar, hover over one end of it until the cursor changes shape. Then drag the widget to make the scrub bar larger or smaller. To scroll through the data, click and hold the scrub bar and slide it toward the end that you want to see.

Focusing on visual elements

When viewing visuals, you can choose data that you want to focus on or exclude. To perform this choice, choose an element such as a bar or bubble, or a row or column header.

Focusing on or excluding data causes Quick Sight to create a filter and show only the data that you selected.

To remove the filter, choose **Filters** at left and then disable or delete the filter. You can also use **Undo** to remove a filter.

If your visual has a legend that shows categories (dimensions), you can click on the values in the legend to see a menu of available actions. For example, suppose that your bar chart has a field in the **Color** or **Group/Color** field well. The bar chart menu displays the actions that you can choose by clicking or right-clicking on a bar, such as the following:

- Focusing on, or excluding, visual elements
- Changing colors of visual elements
- Drilling down into a hierarchy
- Custom actions activated from the menu, including filtering or URL actions

Excluding visual elements

When viewing visuals, you can choose an element on the visual, and then choose to focus on the element. Elements to focus on can include, for example, a bar or bubble, or a row or column header in the case of a pivot table. The exception is that you can't exclude elements that are mapped to date fields. You can exclude multiple elements on a single chart.

Excluding the element creates a filter that removes only that element from the visual.

To see the excluded element again, you can either choose **Undo** on the application bar, or you can disable or delete the filter.

For more information about filters, see [Filtering data in Amazon Quick Sight](#).

Searching for specific values in your data in Quick Sight

When filtering your visual data, previewing anomalies, or using list or dropdown controls in a dashboard, you can quickly search for values that interest you.

You can search for specific values or all values that contain a specific search query. For example, searching for *al* in a list of U.S. states returns **Alabama**, **Alaska**, and **California**.

You can also use wildcard search to search for all values that match a specific character pattern. For example, you can search for all U.S. states that end with the letters *ia* and narrow the results down to California, Georgia, Pennsylvania, Virginia, and West Virginia.

To search for values in a filter or control, enter a search query in the search bar.

Using wildcard search

The following wildcard characters can be used to find values in Quick Sight filters, list and dropdown controls, and anomaly previews.

- ***** - Use an asterisk symbol to search for values that match zero to many characters in a specific position.
- **?** - Use a question mark to match a single character in a specific position.
- **** - Use a backslash to escape the *****, **?**, or **** wildcard characters and search for them in your query. For example, you can search for phrases that end with a question mark.

Following are examples of how supported wildcard characters can be used in a Quick Sight search query.

- **al** - This query searches for all values with **al** and returns Alabama, Alaska, and California.
- **al*** - This query searches for all values that begin with **al** and end with zero to multiple characters. It returns Alabama, and Alaska in a list of U.S. states.
- ***ia** - This query searches for all values that begin with zero to multiple characters and end with letters **ia**. It returns California, Georgia, Pennsylvania, Virginia, and West Virginia.
- ***al*** - This query searches for all values with zero to multiple characters before and after the letters **al**. It returns Alabama, Alaska, and California.
- **a?a?a?a** - This query searches for all values with a single character in the exact positions between the **a** letters. It returns Alabama.
- **a?a*a** - This query searches for all values with a single character between the first two **a** letters and multiple characters between the second two **a** letters. It returns Alabama and Alaska.
- **How*\?** - This query searches for values that begin with **How**, followed by zero to multiple characters, and end with a question mark. The backslash (****) in this query informs Quick Sight to

search for question marks in each value, rather than use the question mark symbol as a wildcard character. This query returns the questions, How are you? and, How is this possible?

- `\**` - This query searches for values that begin with an asterisk and are followed by zero to multiple characters. The backslash (\) in this query informs Quick Sight to search for an actual asterisk in the values, rather than use the asterisk symbol as a wildcard character. This query returns values such as `*all`, `*above`, and `*below`.
- `\\*` - This query searches for values with a backslash, followed by zero to multiple characters. The first backslash (\) in this query informs Quick Sight to search for the second backslash (\) in each value, rather than use the backslash symbol as a wildcard character. This query returns results such as `\Home`.
- `???` - This query searches for values that contain three characters. It returns values such as `ant`, `bug`, and `car`.

Exporting data from visuals

Note

Export files can directly return information from the dataset import. This makes the files vulnerable to CSV injection if the imported data contains formulas or commands. For this reason, export files can prompt security warnings. To avoid malicious activity, turn off links and macros when reading exported files.

Using the Amazon Quick console, you can export data from any type of chart or graph. The export contains only the data in the fields that are currently visible in the selected visualization. Any data that is filtered out is excluded from the export file. You can export data into the following formats:

- A text file containing comma-separated values (CSV), available for all visual types.
- A Microsoft Excel workbook file (.xlsx), available for pivot tables and table charts only.

The following rules apply:

- Exported files are downloaded to the default download directory configured in the browser that you're currently using.
- The downloaded file is named for the visualization that you exported it from. To make the file name unique, it has a sequential timestamp (a Unix epoch data type).

- Default limit for export to CSV format: 500 MB or 1M rows whichever comes first
- Default limit for export to Excel format:
 - from Pivot Table visual 400K cells or 50K rows
 - from Table visual 800K cells or 100K rows

 Note

With a subscription to Paginated Reporting, you are able to [schedule the export of visuals in CSV and Excel formats](#) and export up to 3M rows (CSV) and 16M cells (Excel).

- You can't export data from an insight, because insights consume the data, but don't contain the data.
- Quick Sight doesn't support exporting data from more than a single visualization at a time. To export data from additional visuals in the same analysis or dashboard, repeat this process for each visual. To export all the data from a dashboard or analysis, you need to connect to the original data source using valid credentials and a tool that you can use to extract data.

Use the following procedure to export data from a visualization in Amazon Quick Sight. Before you begin, open the analysis or dashboard that contains the data that you want to export.

To export data from a visualization

1. Choose the visualization that you want to export. Make sure that it is selected and highlighted.
2. At top right on the visual, open the menu and choose one of the following:
 - To export to CSV, choose **Export to CSV**.
 - To export to XSLX, choose **Export to Excel**. This option is available only for pivot tables and table charts.
3. Depending on your browser settings, one of the following happens:
 - The file automatically goes to your default **Download** location.
 - A dialog box appears so you can choose a file name and location.
 - A dialog box appears so you can choose to open the file with the default software or to save to.

Refreshing visuals in Quick Sight

When you work in an Quick Sight analysis or dashboard, visuals refresh and reload when you change something that affects them, such as updating a parameter or filter control. If you switch to a new sheet after a parameter or filter changes, only the visuals affected by the change refresh on the new sheet. Otherwise, visuals update every 30 minutes when you switch sheets. This is the default behavior for all analyses and dashboards.

If you want to refresh all visuals when you switch sheets, regardless of a change, you can do so for each analysis that you create.

To refresh all visuals each time that you switch sheets in an analysis

1. In Amazon Quick, open the analysis.
2. In the analysis, choose **Edit > Analysis Settings**.
3. In the **Analysis Settings** pane that opens, for **Refresh Options**, toggle on **Reload visuals each time I switch sheets**.
4. Choose **Apply**.

Deleting Amazon Quick Sight visuals

Use the following procedure to delete a visual.

To delete a visual

1. On the analysis page, choose the visual that you want to delete.
2. Choose the on-visual menu at the upper-right corner of the visual, and then choose **Delete**.

Using Topics on sheets in Amazon Quick Sight

You can use Quick Sight Topics as the data model for your analysis sheets. When you create an analysis from a Topic, you can select fields from any of the datasets in the Topic and Quick Sight automatically performs runtime inner joins based on the defined relationships. This eliminates the need to pre-join tables into a single flat dataset before building an analysis.

Creating an analysis from a Topic

To use a Topic as the data model for an analysis, use the following procedure.

To create an analysis from a Topic

1. On the Quick homepage, choose **Data** in the navigation pane, then choose the **Topics** tab.
2. Open the Topic you want to use.
3. Choose **Create analysis**.
4. Select **Interactive sheet** and choose **Create**.
5. In the analysis, the fields list shows columns from all datasets in the Topic. Select fields from multiple datasets and add them to a visual. Quick Sight performs runtime inner joins automatically based on the defined relationships.

The following rules apply when working with Topics in analysis sheets:

- You must have Owner or Viewer permissions on the Topic to use it in an analysis.
- When a visual references fields from multiple datasets, Quick Sight uses the defined relationships to determine the join path and performs inner joins at runtime.
- You can create calculated fields that reference columns from different datasets within the Topic.
- Row-level security (RLS) rules are enforced during runtime joins. Users see only the intersection of rows they are permitted to access in each dataset.
- Each dataset in the Topic can have its own independent refresh schedule. Visuals always reflect the latest data available in each dataset.

Note

Runtime joins in analysis sheets use inner join semantics only. Rows without matching keys in both datasets are excluded from visual results. If you need outer join behavior, use the Topic in Amazon Quick chat, where the LLM-powered chat agent can generate SQL with outer joins based on custom instructions.

For more information about creating and configuring Topics, see [Working with Amazon Quick Sight Topics](#).

Enabling legacy Topics in an analysis

Note

The following applies to legacy Topics only. For information about the differences between new Topics and legacy Topics, see [Working with legacy Topics](#).

For legacy Topics, you can enable a topic in an analysis to activate the ML-powered automated data prep, which speeds natural language topic creation. Automated data prep automatically selects high-value fields, chooses user-friendly field names and synonyms, and formats data for presentation.

Automated data prep binds the legacy Topic to your analysis and prepares an index for searching in natural language. Dashboard users find that the linked topic is automatically selected in the search bar, making it easier for them to query the dataset.

The following rules apply to working with legacy Topics in an analysis:

- You must be an owner of the underlying dataset before you can create a legacy Topic using that dataset or an analysis that uses that dataset.
- You must be an owner of a legacy Topic before you can link it to an analysis.

To enable a legacy Topic

1. Open the analysis that you want to use with automated data prep.
2. On the top navigation bar, choose the topic icon.
3. Choose one of the following:
 - To activate a new topic, select **Create new topic** and enter a topic title and optional description.
 - To activate an existing topic, select **Update existing topic** and choose the topic from the list.
4. Choose **ENABLE TOPIC** to confirm your choice.
5. When the topic is finished processing, you can use what it learned from the analysis to ask questions in natural language.

Now, when users navigate to the dashboard, the linked topic is automatically selected in the search bar.

After a legacy Topic is linked to an analysis, further updates to the analysis are not automatically synced to the topic. Authors need to manage updating topics manually from the **Topics** page.

Visual types in Amazon Quick Sight

Amazon Quick Sight offers a range of visual types that you can use to display your data. Use the topics in this section to learn more about the capabilities of each visual type.

Topics

- [Measures and dimensions](#)
- [Display limits](#)
- [Hiding or displaying the other category](#)
- [Customizing the number of data points to display](#)
- [Using AutoGraph](#)
- [Using bar charts](#)
- [Using box plots](#)
- [Using combo charts](#)
- [Using custom visual content](#)
- [Using donut charts](#)
- [Using funnel charts](#)
- [Using gauge charts](#)
- [Using heat maps](#)
- [Using Highcharts](#)
- [Using histograms](#)
- [Using image components](#)
- [Using KPIs](#)
- [Using layer maps](#)
- [Using line charts](#)
- [Creating maps and geospatial charts](#)
- [Using small multiples](#)
- [Using pie charts](#)

- [Using pivot tables](#)
- [Using radar charts](#)
- [Using Sankey diagrams](#)
- [Using scatter plots](#)
- [Using tables as visuals](#)
- [Using text boxes](#)
- [Using tree maps](#)
- [Using waterfall charts](#)
- [Using word clouds](#)

Measures and dimensions

We use the term *measure* to refer to numeric values that you use for measurement, comparison, and aggregation in visuals. A measure can be either a numeric field, like product cost, or a numeric aggregate on a field of any data type, like count of transaction IDs.

We use the term *dimension* or *category* to refer to text or date fields that can be items, like products, or attributes that are related to measures and can be used to partition them. Examples are sales date for sales figures or product manufacturer for customer satisfaction numbers. Amazon Quick Sight automatically identifies a field as a measure or a dimension based on its data type.

Numeric fields can act as dimensions, for example ZIP codes and most ID numbers. It's helpful to give such fields a string data type during data preparation. This way, Amazon Quick Sight understands that they are to be treated as dimensions and are not useful for performing mathematical calculations.

You can change whether a field is displayed as a dimension or measure on an analysis-by-analysis basis instead. For more information, see [Fields as dimensions and measures](#).

Display limits

All visual types limit the number of data points they display, so that the visual elements (like lines, bars, or bubbles) are still easy to view and analyze. The visual selects the first n number of rows for display up to the limit for that visual type. The selection is either according to sort order, if one has been applied, or in default order otherwise.

The number of data points supported varies by visual type. To learn more about display limits for a particular visual type, see the topic for that type.

The visual title identifies the number of data points displayed if you have reached the display limit for that visual type. If you have a large dataset and want to avoid running into the visual display limit, use one or more filters to reduce the amount of data displayed. For more information about using filters with visuals, see [Filtering data in Amazon Quick Sight](#).

For dashboards and analyses, Amazon Quick Sight supports the following:

- 50 datasets per dashboard
- 20 sheets per dashboard
- 30 visualization objects per sheet

Note

Amazon Quick Sight supports over 30 different *visual types* (categories of charts and visualizations such as bar charts, pie charts, and line charts). Each analysis sheet can contain up to 30 *visual instances* (individual chart objects) of any combination of types.

You can also choose to limit how many data points you want to display in your visual, before they are added to the **other** category. This category contains the aggregated data for all the data beyond the cutoff limit for the visual type you are using—either the one you impose, or the one based on display limits. You can use the on-visual menu to choose whether to display the **other** category. The **other** category doesn't show on scatter plots, heat maps, maps, tables (tabular reports), or key performance indicators (KPIs). It also doesn't show on line charts when the x-axis is a date. Drilling down into the **other** category is not supported.

Hiding or displaying the other category

Use the following procedure to hide or display the "other" category.

To hide or display the "other" category

1. On the analysis page, choose the visual that you want to modify.
2. Choose the on-visual menu at the upper-right corner of the visual, and then choose **Hide "other" category** or **Show "other" category**, as appropriate.

Customizing the number of data points to display

You can choose the number of data points to display on the main axis of some visuals. After this number is displayed in the chart, any additional data points are included in the "other" category. For example, if you choose to include 10 data points out of 200, 10 display in the chart and 190 become part of the "other" category.

To find this setting, choose the v-shaped on-visual menu, then choose **Format visual**. You can use the following table to determine which field well contains the data point setting and what number of data points the visual type displays by default.

Visual type	Where to find the data point setting	Default number of data points
Bar chart, horizontal	Y-axis – Number of data points displayed	10,000
Bar chart, vertical	X-axis – Number of data points displayed	10,000
Combo chart	X-axis – Number of data points displayed	2,500
Heat map	Rows – Number of rows displayed Columns – Number of columns displayed	100
Line chart	X-axis – Number of data points displayed	10,000
Pie chart	Group/Color – Number of slices displayed	20
Tree map	Group by – Number of squares displayed	100

Using AutoGraph

AutoGraph isn't a visual type itself, but instead lets you tell Amazon Quick to choose the visual type for you. When you create a visual by choosing AutoGraph and then selecting fields, Amazon Quick uses the most appropriate visual type for the number and data types of the fields you select.

Creating a visual using AutoGraph

Use the following procedure to create a visual using AutoGraph.

To create a visual using AutoGraph

1. On the analysis page, choose **Visualize** on the tool bar.
2. Choose **Add** on the application bar, and then choose **Add visual**.
3. On the **Visual types** pane, choose the AutoGraph icon.
4. On the **Fields list** pane, choose the fields that you want to use.

Using bar charts

Amazon Quick supports the following types of bar charts, with either horizontal or vertical orientation:

- **Single-measure** – A *single-measure bar chart* shows values for a single measure for a dimension.
- **Multi-measure** – A *multi-measure bar chart* shows values for multiple measure for a dimension.
- **Clustered** – A *clustered bar chart* shows values for a single measure for a dimension, grouped by another dimension.
- **Stacked** – A *stacked bar chart* is similar to a clustered bar chart in that it displays a measure for two dimensions. However, instead of clustering bars for each child dimension by the parent dimension, it displays one bar per parent dimension. It uses color blocks within the bars to show the relative values of each item in the child dimension. The color blocks reflect the value of each item in the child dimension relative to the total for the measure. A stacked bar chart uses a scale based on the maximum value for the selected measure.
- **Stacked 100 percent** – A *stacked 100 percent bar chart* is similar to a stacked bar chart. However, in a stacked 100 percent bar chart, the color blocks reflect the percentage of each item in the child dimension, out of 100 percent.

Bar charts show up to 10,000 data points on the axis for visuals that don't use group or color. For visuals that do use group or color, they show up to 50 data points on the axis and up to 50 data points for group or color. For more information about how Amazon Quick handles data that falls outside display limits, see [Display limits](#).

Creating single-measure bar charts

Use the following procedure to create a single-measure bar chart.

To create a single-measure bar chart

1. On the analysis page, choose **Visualize** on the toolbar at left.
2. On the application bar at upper left, choose **Add**, and then choose **Add visual**.
3. On the **Visual types** pane, choose the **Horizontal bar chart** or **Vertical bar chart** icon.
4. From the **Fields list** pane, drag a dimension to the **X-axis** or **Y-axis** field well.
5. From the **Fields list** pane, drag a measure to the **Value** field well.

Creating multi-measure bar charts

Use the following procedure to create a multi-measure bar chart.

To create a multi-measure bar chart

1. On the analysis page, choose **Visualize** on the toolbar at left.
2. On the application bar at upper-left, choose **Add**, and then choose **Add visual**.
3. On the **Visual types** pane, choose the **Horizontal bar chart** or **Vertical bar chart** icon.
4. From the **Fields list** pane, drag a dimension to the **X-axis** or **Y-axis** field well.
5. From the **Fields list** pane, drag two or more measures to the **Value** field well.

Creating clustered bar charts

Use the following procedure to create a clustered bar chart.

To create a clustered bar chart

1. On the analysis page, choose **Visualize** on the toolbar at left.
2. On the application bar at upper left, choose **Add**, and then choose **Add visual**.
3. On the **Visual types** pane, choose the **Horizontal bar chart** or **Vertical bar chart** icon.
4. From the **Fields list** pane, drag a dimension to the **X-axis** or **Y-axis** field well.
5. From the **Fields list** pane, drag a measure to the **Value** field well.
6. From the **Fields list** pane, drag a dimension to the **Group/Color** field well.

Creating stacked bar charts

Use the following procedure to create a stacked bar chart.

To create a stacked bar chart

1. On the analysis page, choose **Visualize** on the toolbar at left.
2. On the application bar at upper-left, choose **Add**, and then choose **Add visual**.
3. On the **Visual types** pane, choose the **Horizontal stacked bar chart** or **Vertical stacked bar chart** icon.
4. From the **Fields list** pane, drag a dimension to the **X-axis** or **Y-axis** field well.
5. From the **Fields list** pane, drag a dimension to the **Group/Color** field well.
6. From the **Fields list** pane, drag a measure to the **Value** field well.
7. (Optional) Add data labels and show totals:
 - a. On the menu in the upper-right corner of the visual, choose the **Format visual** icon.
 - b. In the **Visual** pane, choose **Data labels**.
 - c. Toggle the switch to display data labels.

Labels for each measure value appear in the chart and the option to show totals appears in the pane.

- d. Check **Show totals**.

Totals appear for each bar in the chart.

Creating stacked 100 percent bar charts

Use the following procedure to create a stacked 100 percent bar chart.

To create a stacked 100 percent bar chart

1. On the analysis page, choose **Visualize** on the toolbar at left.
2. On the application bar at upper-left, choose **Add**, and then choose **Add visual**.
3. On the **Visual types** pane, choose the **Horizontal stacked 100% bar chart** or **Vertical stacked 100% bar chart** icon.
4. From the **Fields list** pane, drag a dimension to the **X-axis** or **Y-axis** field well.
5. From the **Fields list** pane, drag two or more measures to the **Value** field well.

Bar chart features

To understand the features supported by bar charts, use the following table.

Feature	Supported?	Comments	For more information
Changing the legend display	Yes, with exceptions	Multi-measure and clustered bar charts display a legend, while single-measure horizontal bar charts don't.	Legends on visual types in Quick
Changing the title display	Yes		Titles and subtitles on visual types in Quick
Changing the axis range	Yes		Range and scale on visual types in Quick
Showing or hiding axis lines, grid lines, axis labels, and axis sort icons	Yes		Axes and grid lines on visual types in Quick
Changing the visual colors	Yes		Colors in visual types in Quick
Focusing on or excluding elements	Yes, with exceptions	You can focus on or exclude any bar on the chart, except when you are using a date field as the dimension for the axis. In that case, you can only focus on a bar, not exclude it.	Focusing on visual elements Excluding visual elements
Sorting	Yes	You can sort on the fields you choose for the axis and the values.	Sorting visual data in Amazon Quick

Feature	Supported?	Comments	For more information
Performing field aggregation	Yes	You must apply aggregation to the field or fields you choose for the value, and can't apply aggregation to the fields you choose for the axis or group/ color.	Changing field aggregation
Adding drill-downs	Yes	You can add drill-down levels to the axis and Group/Color field wells.	Adding drill-downs to visual data in Quick Sight
Showing data labels	Yes		Data labels on visual types in Quick
Showing stacked bar chart totals	Yes	Showing totals in a stacked bar chart is only available when you choose to show data labels.	Stacked bar charts

Using box plots

Box plots, also known as box and whisker plots, display data pooled from multiple sources into one visual, helping you make data-driven decisions. Use a box plot to visualize how data is distributed across an axis or over time, for example flights delayed over a 7-day time period. Typically, a box plot details information in quarters:

- **Minimum** – The lowest data point excluding outliers.
- **Maximum** – The highest data point excluding outliers.
- **Median** – The middle value of the dataset.
- **First Quartile** – The middle value between the smallest number and the median of the dataset. The first quartile doesn't include the minimum or the median.
- **Third Quartile** – The middle value between the largest number and the median of the dataset. The third quartile doesn't include the maximum or the median.

Outliers are extreme data points that aren't included in the calculation of a box plot's key values. Because outliers are calculated separately, their data points don't appear immediately after a box plot is created. Box plots display up to 10,000 data points. If a dataset contains more than 10,000 data points, a warning appears at the upper-right corner of the visual.

Box plots support up to five metrics and one group-by, but don't render if duplicate metrics are supplied.

Box plots support some calculated fields, but not all. Any calculated field that uses a window function, for example `avgOver1`, results in a SQL error.

Box plot visuals aren't compatible with MySQL 5.3 and earlier.

To create a basic box plot visual

1. Sign in to Amazon Quick at <https://quicksight.aws.amazon.com/>.
2. Open Quick and choose **Analyses** on the navigation pane at left.
3. Choose one of the following:
 - To create a new analysis, choose **New analysis** at upper right. For more information, see [Starting an analysis in Quick Sight](#).
 - To use an existing analysis, choose the analyses that you want to edit.
4. Choose **Add, Add visual**.
5. At lower left, choose the box plot icon from **Visual types**.
6. On the **Fields list** pane, choose the fields that you want to use for the appropriate field wells. Box plots require at least one unique measure field.
7. (Optional) Add drill-down layers by dragging one or more additional fields to the **Group/Color** field well. For more information about adding drill-downs, see [Adding drill-downs to visual data in Quick Sight](#).

To understand the features supported by box plots, see [Analytics formatting per type in Quick](#). For customization options, see [Formatting in Amazon Quick](#).

Using combo charts

Using a combo chart, you can create one visualization that shows two different types of data, for example trends and categories. Combo charts are also known as line and column (bar) charts, because they combine a line chart with a bar chart. Bar charts are useful for comparing categories.

Both bar charts and line charts are useful for displaying changes over time, although bar charts should show a greater difference between changes.

Amazon Quick supports the following types of combo charts:

- **Clustered bar combo charts** – display sets of single-color bars where each set represents a parent dimension and each bar represents a child dimension. Use this chart to make it easy to determine values for each bar.
- **Stacked bar combo charts** – display multi-color bars where each bar represents a parent dimension and each color represents a child dimension. Use this chart to make it easy to see relationships between child dimensions within a parent dimension. This chart shows the total value for the parent dimension and how each child adds to the total value. To determine the value for each child dimension, the chart reader must compare the size of the color section to the data labels for that axis.

Both types of combo chart require only one dimension on the **X axis**, but are usually more effective when also displaying at least one measure under **Lines**.

Use a combo chart only if you want to show a relationship between the bars and the lines. A good rule of thumb is that if you need to explain how the two chart types relate, you should probably use two separate charts instead.

Because each chart works differently, it can be helpful to understand the following points before you begin:

- The data points in each series render on different scales. Combo charts use a scale based on the maximum value for the selected measure.
- The distance between the numbers on the axis won't match between the lines and bars, even if you select the same scale for each chart type.
- For clarity, try to use different units for the measure in each data series.

The combo chart is like using two different types of visualization at the same time. Make sure that the data in the bars (or columns) directly relates to the data in the line or lines. This relationship is not technically enforced by the tool, so it's essential that you determine this relationship yourself. Without some relation between the lines and bars, the visual loses meaning.

You can use the combo chart visual type to create a single-measure or single-line chart. A single-measure combo chart shows one measure for one dimension.

To create a multi-measure chart, you can choose to add multiple lines, or multiple bars. A multi-measure bar chart shows two or more measures for one dimension. You can group the bars in clusters, or stack them.

For the bars, use a dimension for the axis and a measure for the value. The dimension is typically a text field that is related to the measure in some way and can be used to segment it to see more detailed information. Each bar in the chart represents a measure value for an item in the dimension you chose.

Bars and lines show up to 2,500 data points on the axis for visuals that don't use group or color. For visuals that do use group or color, bars show up to 50 data points on the axis and up to 50 data points for group or color, while lines show 200 data points on the axis and up to 25 data points for group or color. For more information about how Amazon Quick handles data that falls outside display limits, see [Display limits](#).

Combo chart features

To understand the features supported by combo charts, use the following table.

Feature	Supported?	Comments	For more information
Changing the legend display	Yes, with exceptions	Multi-measure combo charts display a legend, and single-measure combo charts don't.	Legends on visual types in Quick
Changing the title display	Yes		Titles and subtitles on visual types in Quick
Changing the axis range	Yes	You can set the range for the axis.	Range and scale on visual types in Quick
Showing or hiding axis lines, grid lines, axis labels, and axis sort icons	Yes		Axes and grid lines on visual types in Quick


Feature	Supported?	Comments	For more information
Changing the visual colors	Yes		Colors in visual types in Quick
Focusing on or excluding elements	Yes, with exceptions	You can focus on or exclude any bar on the chart, except when you are using a date field as the dimension for the axis. In that case, you can only focus on a bar, not exclude it.	Focusing on visual elements Excluding visual elements
Sorting	Yes	You can sort on the fields you choose for the axis and the values.	Sorting visual data in Amazon Quick
Performing field aggregation	Yes	You must apply aggregation to the field or fields you choose for the value. You can't apply aggregation to the fields you choose for the axis or group/color.	Changing field aggregation
Adding drill-downs	Yes	You can add drill-down levels to the axis and Group/Color field wells.	Adding drill-downs to visual data in Quick Sight
Synchronizing y-axis	Yes	Synchronize the y-axes for both bars and lines into a single axis.	Range and scale on visual types in Quick

Creating a combo chart

Use the following procedure to create a combo chart.

To create a combo chart

1. On the analysis page, choose **Visualize** on the tool bar.
2. Choose **Add** on the application bar, and then choose **Add visual**.
3. On the **Visual types** pane, choose one of the combo chart icons.
4. From the **Fields list** pane, drag the fields that you want to use to the appropriate field wells. Typically, you want to use dimension or measure fields as indicated by the target field well. If you choose to use a dimension field as a measure, the **Count** aggregate function is automatically applied to it to create a numeric value. You can create combo charts as follows:
 - Choose a dimension for the **X axis**.
 - To create a single-measure combo chart, choose one measure for either **Bars** or **Lines**.
 - To create a multi-measure combo chart, choose two or more measures for the **Bars** or **Lines** field well.
 - Optionally, add a dimension to the **Group/Color** field well. If you have a field in **Group/Color**, you can't have more than one field under **Bars**.



5. (Optional) Add drill-down layers by dragging one or more additional fields to the **X axis** or **Group/Color** field wells. For more information about adding drill-downs, see [Adding drill-downs to visual data in Quick Sight](#).

Using custom visual content

You can embed webpages and online videos, forms, and images in your Quick dashboards using the custom visual content chart type.

For example, you can embed the image of your company logo in your dashboards. You can also embed an online video from your organization's latest conference, or embed an online form asking readers of the dashboard if the dashboard is helpful.

After you create custom visual content, you can use navigation actions to navigate within them. You can also use parameters to control what appears in them.

The following limitations apply to custom visual content:

- Only `https` URL schemes are supported.
- Custom visual content isn't supported in email reports.
- Images and websites that use hotlink protection won't load in custom visuals.

To embed a webpage, video, online form, or image in your dashboard, choose the custom visual content icon in the **Visual types** pane.

For more information about adding visuals to a dashboard, see [Adding a visual](#).

Use the following procedures to learn how to embed custom visuals in your dashboards.

Best practices for using custom visual content

When embedding web content using the custom visual content chart type, we recommend the following:

- Choose web content from sources that support viewing or opening the content in an `IFrame`. If the source of the web content doesn't support being viewed or opened in an `IFrame`, the content doesn't appear in Quick, even if the URL is accurate.
- When possible, use embeddable URLs, especially for videos, online forms, spreadsheets, and documents. Embeddable URLs create a better experience for readers of your dashboard and make interacting with the content easier. You can usually find the embeddable URL for content when you choose to share the content from the source website.
- To embed internal URLs or URLs that you own, you might need to set them to be opened in an `IFrame`.

- When viewing custom visual content in an analysis or dashboard, make sure you enable all cookies. If third-party cookies are blocked in your browser, images that are part of the website that is embedded within the custom content visual do not render.

Note

Chrome has announced plans to deprecate all third-party cookies by the end of 2024. This means that websites that are embedded within Quick custom content visuals will no longer show any contents that rely on third-party cookies in Chrome. For more information about Chrome's plans to deprecate third party cookies, see [Chrome is deprecating third-party cookies](#).

Embedding images in a dashboard

You can embed an online image in a dashboard using the image URL. Use the following procedure to embed an image using the custom visual content chart type.

Embedded images don't appear in a browser that has third-party cookies blocked. To see embedded images in a dashboard, enable third-party cookies in your browser settings.

To embed an image in a dashboard

1. In the **Visual types** pane, choose the custom visual content icon.
2. In the visual, choose **Customize visual**.
3. In the **Properties** pane that opens, under **Custom content**, enter the image URL for the image that you want to embed.
4. Choose **Apply**.

The image appears as a webpage in the visual.

5. Choose **Show as image**.

If the URL is an image, the image appears in the visual.

If the URL is not an image, such as a URL to a slide show, gallery, or webpage, the following message appears: This URL doesn't appear to be an image. Update the URL to an image. To do so, open the image that you want to embed in a separate browser tab, or choose an embeddable URL for the image (usually found when you choose to share the image).

6. (Optional) For **Image sizing options**, choose one of the following options:

- **Fit to width** – This option fits the image to the width of the visual.
- **Fit to height** – This option fits the image to the height of the visual.
- **Scale to visual** – This option scales the image to the width and height of the visual. This option might contort the image.
- **Do not scale** – This option keeps the image at its original scale and doesn't fit the image to the dimensions of the visual. With this option, the image is centered in the visual and the parts of the image that are within the width and height of the visual appear. Some parts of the image might not appear if the visual is smaller than the image. If the visual is larger than the image, however, the image is centered in the visual and is surrounded by white space.

Embedding online forms in a dashboard

You can embed an online form in a dashboard using the embeddable URL. Use the following procedure to embed an online form using the custom visual content chart type.

To embed an online form in a dashboard

1. In the **Visual types** pane, choose the custom visual content icon.
2. In the visual, choose **Customize visual**.
3. In the **Properties** pane that opens, under **Custom content**, enter the form URL for the online form that you want to embed.

If possible, use an embeddable URL for the form. Using an embeddable URL creates a better experience for readers of your dashboard who might want to interact with the form. You can often find the embeddable URL when you choose to share the form on the site where you create it.

4. Choose **Apply**.

The form appears in the visual.

Embedding webpages in a dashboard

You can embed webpage in a dashboard using the URL. Use the following procedure to embed webpage using the custom visual content chart type.

To embed a webpage in a dashboard

1. In the **Visual types** pane, choose the custom visual content icon.
2. In the visual, choose **Customize visual**.
3. In the **Properties** pane that opens, under **Custom content**, enter the URL for the webpage that you want to embed.
4. Choose **Apply**.

The webpage appears in the visual.

Embedding online videos in a dashboard

You can embed an online video in a dashboard using the embeddable video URL. Use the following procedure to embed an online video using the custom visual content chart type.

To embed an online video in a dashboard

1. In the **Visual types** pane, choose the custom visual content icon.
2. In the visual, choose **Customize visual**.
3. In the **Properties** pane that opens, under **Custom content**, enter the embeddable URL for the video that you want to embed.

To find the embeddable URL for a video, share the video and copy the embed URL from IFrame code. The following is an example of an embed URL for a YouTube video: `https://www.youtube.com/embed/uniqueid`. For a Vimeo video, the following is an example of an embed URL: `https://player.vimeo.com/video/uniqueid`.

4. Choose **Apply**.

The video appears in the visual.

Using donut charts

Use donut charts to compare values for items in a dimension. The best use for this type of chart is to show a percentage of a total amount.

Each wedge in a donut chart represents one value in a dimension. The size of the wedge represents the proportion of the value for the selected measure that the item represents compared to the

whole for the dimension. Donut charts are best when precision isn't important and there are few items in the dimension.

To learn how to use donut charts in Amazon Quick, you can watch this video:

[Using Donut Charts](#)

To create a donut chart, use one dimension in the **Group/Color** field well. With only one field, the chart displays the division of values by row count. To display the division of dimension values by a metric value, you can add a metric field to the **Value** field well.

Donut charts show up to 20 data points for group or color. For more information about how Amazon Quick handles data that falls outside display limits, see [Display limits](#).

Donut chart features

To understand the features supported by donut charts, use the following table.

Feature	Supported?	Comments	For more information
Changing the legend display	Yes		Legends on visual types in Quick
Changing the title display	Yes		Titles and subtitles on visual types in Quick
Changing the axis range	Not applicable		Range and scale on visual types in Quick
Changing the visual colors	Yes		Colors in visual types in Quick
Focusing on or excluding elements	Yes, with exceptions	You can focus on or exclude a wedge in a donut chart, except when you are using a date field as a dimension . In that case, you can only focus on a wedge, not exclude it.	Focusing on visual elements Excluding visual elements

Feature	Supported?	Comments	For more information
Sorting	Yes	You can sort on the field that you choose for the value or the group or color.	Sorting visual data in Amazon Quick
Performing field aggregation	Yes	You must apply aggregation to the field that you choose for the value, and can't apply aggregation to the field that you choose for group or color.	Changing field aggregation
Adding drill-downs	Yes	You can add drill-down levels to the Group/Color field well.	Adding drill-downs to visual data in Quick Sight
Choosing size	Yes	You can choose how thick the donut chart is: small, medium, and large.	Formatting in Amazon Quick
Showing totals	Yes	You can choose to display or hide the aggregate of the Value field. By default, this displays the total count of the Group/Color field, or the total sum of the Value field.	Formatting in Amazon Quick

Creating a donut chart

Use the following procedure to create a donut chart.

To create a donut chart

1. On the analysis page, choose **Visualize** on the tool bar.
2. Choose **Add** on the application bar, and then choose **Add visual**.
3. On the **Visual types** pane, choose the donut chart icon.

4. From the **Fields list** pane, drag the fields that you want to use to the appropriate field wells. Typically, you want to use dimension or measure fields as indicated by the target field well. If you choose to use a dimension field as a measure, the **Count** aggregate function is automatically applied to it to create a numeric value.

To create a donut chart, drag a dimension to the **Group/Color** field well. Optionally, drag a measure to the **Value** field well.

5. (Optional) Add drill-down layers by dragging one or more additional fields to the **Group/Color** field well. For more information about adding drill-downs, see [Adding drill-downs to visual data in Quick Sight](#).

Using funnel charts

Use a funnel chart to visualize data that moves across multiple stages in a linear process. In a funnel chart, each stage of a process is represented in blocks of different shapes and colors. The first stage, known as the *head*, is the largest block and is followed by the smaller stages, known as the *neck*, in a funnel shape. The size of the block representing each stage in a funnel chart is a percentage of the total, and is proportionate to its value. The bigger the size of the block, the bigger its value.

Funnel charts are often useful in business contexts because you can view trends or potential problem areas in each stage, such as bottlenecks. For example, they can help you visualize the amount of the potential revenue in each stage of a sale, from first contact to final sale and on through maintenance.

To create a basic funnel chart visual

1. Open Amazon Quick and choose **Analyses** on the navigation pane at left.
2. Choose one of the following:
 - To create a new analysis, choose **New analysis** at upper right. For more information, see [Starting an analysis in Quick Sight](#).
 - To use an existing analysis, choose the analysis that you want to edit.
3. Choose **Add (+), Add Visual**.
4. At lower left, choose the funnel chart icon from **Visual types**.
5. On the **Fields list** pane, choose the fields that you want to use for the appropriate field wells. Funnel charts require one dimension in **Group**.

6. (Optional) Add drill-down layers by dragging one or more additional fields to the **Group/Color** field well. For more information about adding drill-downs, see [Adding drill-downs to visual data in Quick Sight](#).

To understand the features supported by funnel charts, see [Analytics formatting per type in Quick](#). For customization options, see [Formatting in Amazon Quick](#).

Using gauge charts

Use gauge charts to compare values for items in a measure. You can compare them to another measure or to a custom amount.

A gauge chart is similar to a nondigital gauge, for example a gas gauge in an automobile. It displays how much there is of the thing you are measuring. In a gauge chart, this measurement can exist alone or in relation to another measurement. Each color section in a gauge chart represents one value. In the following example, we are comparing actual sales to the sales goal, and the gauge shows that we must sell an additional 33.27% to meet the goal.

To learn how to use gauge charts in Amazon Quick, you can watch this video:

[Using Gauge Charts](#)

To create a gauge chart, you need to use at least one measure. Put the measure in the **Value** field well. If you want to compare two measures, put the additional measure in the **Target value** field well. If you want to compare a single measure to a target value that isn't in your dataset, you can use a calculated field that contains a fixed value.

You can choose a variety of formatting options for the gauge chart, including the following settings in **Format visual**.

- **Value displayed** – Hide value, display actual value, or display a comparison of two values
- **Comparison method** – Compare values as a percent, the actual difference between values, or difference as a percent
- **Axis style** –
 - **Show axis label** – Show or hide the axis label
 - **Range** – The numeric minimum and maximum range to display in the gauge chart
 - **Reserve padding (%)** – Added to the top of the range (target, actual value, or max)
- **Arc style** – Degrees the arc displays (180° to 360°)

- **Thickness** – Thickness of the arc (small, medial, or large)

Gauge chart features

To understand the features supported by gauge charts, use the following table.

Feature	Supported?	Comments	For more information
Changing the legend display	Yes		Legends on visual types in Quick
Changing the title display	Yes		Titles and subtitles on visual types in Quick
Formatting gauge	Yes	You can customize the value displayed, the comparison method, the axis style, the arc style, and the thickness of the gauge.	
Changing the axis range	No		
Changing the visual colors	Yes	The foreground color the filled area; it represents the Value . The background color the unfilled area; it represents the Target value if one is selected.	Colors in visual types in Quick
Focusing on or excluding elements	No		
Sorting	No		Sorting visual data in Amazon Quick

Feature	Supported?	Comments	For more information
Performing field aggregation	Yes		Changing field aggregation
Adding drill-downs	No		

Creating a gauge chart

Use the following procedure to create a gauge chart.

To create a gauge chart

1. On the analysis page, choose **Visualize** on the tool bar.
2. Choose **Add** on the application bar, and then choose **Add visual**.
3. On the **Visual types** pane, choose the gauge chart icon.
4. From the **Fields list** pane, drag the fields that you want to use to the appropriate field wells. To create a gauge chart, drag a measure to the **Value** field well. To add a comparison value, drag a different measure to the **Target value** field well.

Using heat maps

Use heat maps to show a measure for the intersection of two dimensions, with color-coding to easily differentiate where values fall in the range. Heat maps can also be used to show the count of values for the intersection of the two dimensions.

Each rectangle on a heat map represents the value for the specified measure for the intersection of the selected dimensions. Rectangle color represents where the value falls in the range for the measure, with darker colors indicating higher values and lighter colors indicating lower ones.

Heat maps and pivot tables display data in a similar tabular fashion. Use a heat map if you want to identify trends and outliers, because the use of color makes these easier to spot. Use a pivot table if you want to further analyze data on the visual, for example by changing column sort order or applying aggregate functions across rows or columns.

To create a heat map, choose at least two fields of any data type. Amazon Quick populates the rectangle values with the count of the x-axis value for the intersecting y-axis value. Typically, you choose a measure and two dimensions.

Heat maps show up to 50 data points for rows and up to 50 data points for columns. For more information about how Amazon Quick handles data that falls outside display limits, see [Display limits](#).

Heat map features

To understand the features supported by heat maps, use the following table.

Feature	Supported?	Comments	For more information
Changing the legend display	Yes		Legends on visual types in Quick
Changing the title display	Yes		Titles and subtitles on visual types in Quick
Changing the axis range	Not applicable		Range and scale on visual types in Quick
Changing the visual colors	No		Colors in visual types in Quick
Focusing on or excluding elements	Yes, with exceptions	You can focus on or exclude a rectangle in a heat map, except when you are using a date field as the rows dimension. In that case, you can only focus on a rectangle, not exclude it.	Focusing on visual elements Excluding visual elements
Sorting	Yes	You can sort by the fields you choose for the columns and the values.	Sorting visual data in Amazon Quick

Feature	Supported?	Comments	For more information
Performing field aggregation	Yes	You must apply aggregation to the fields you choose for the value, and can't apply aggregation to the fields you choose for the rows or columns.	Changing field aggregation
Adding drill-downs	Yes	You can add drill-down levels to the Rows and Columns field wells.	Adding drill-downs to visual data in Quick Sight
Conditional formatting	No		Conditional formatting on visual types in Quick

Creating a heat map

Use the following procedure to create a heat map.

To create a heat map

1. On the analysis page, choose **Visualize** on the tool bar.
2. Choose **Add** on the application bar, and then choose **Add visual**.
3. On the **Visual types** pane, choose the heat map icon.
4. From the **Fields list** pane, drag the fields that you want to use to the appropriate field wells. Typically, you want to use dimension or measure fields as indicated by the target field well. If you choose to use a dimension field as a measure, the **Count** aggregate function is automatically applied to it to create a numeric value.

To create a heat map, drag a dimension to the **Rows** field well, a dimension to the **Columns** field well, and a measure to the **Values** field well.

5. (Optional) Add drill-down layers by dragging one or more additional fields to the **Rows** or **Columns** field wells. For more information about adding drill-downs, see [Adding drill-downs to visual data in Quick Sight](#).


Using Highcharts

Use Highcharts visuals to create custom chart types and visuals that use the [Highcharts Core library](#). Highcharts visuals provide Quick authors direct access to the [Highcharts API](#).

To configure a Highcharts visual, Quick authors need to add a Highcharts JSON schema to the visual in Quick. Authors can use Quick expressions to reference Quick fields, and formatting options in the JSON schema that they use to generate the Highcharts visual. The JSON **Chart code** editor provides contextual assistance for autocomplete and real time validation to ensure that the input JSON schemas are configured properly. To maintain security, the Highcharts visual editor does not accept CSS, JavaScript, or HTML code input.

For more information about Highcharts visuals in Amazon Quick, see the [Highcharts Visual QuickStart Guide](#) in [DemoCentral](#).

The following image shows a lipstick chart that is configured in the **Chart code** JSON editor of a Highcharts visual in Quick.



For more examples of visuals that you can create with the Highcharts visual in Quick, see [Highcharts demos](#).

Considerations

Before you start creating Highcharts visuals in Amazon Quick, review the following limitations that apply to Highcharts visuals.

- The following JSON values are not supported in the Highcharts **Chart code** JSON editor:
 - Functions
 - Dates
 - Undefined values
- Links to GeoJSON files or other images are not supported for Highcharts visuals.
- Field colors are not available for Highcharts visuals. Default theme colors are applied to all Highcharts visuals.

Creating a Highcharts visual

Use the following procedure to create a Highcharts visual in Amazon Quick.

1. Open the [Quick console](#).
2. Open the Quick analysis that you want to add a Highcharts visual to.
3. On the application bar, choose **Add**, and then choose **Add visual**.
4. On the **Visual types** pane, choose the Highcharts visual icon. An empty visual appears on the analysis sheet and the **Properties** pane opens on the left.
5. In the **Properties** pane, expand the **Display settings** section and perform the following actions:
 - a. For **Edit title**, choose the paintbrush icon, enter the title that you want the visual to have, and then choose **SAVE**. Alternatively, choose the eyeball icon to hide the title.
 - b. (Optional) For **Edit subtitle**, choose the paintbrush icon, enter the subtitle that you want the visual to have, and then choose **SAVE**. Alternatively, choose the eyeball icon to hide the subtitle.
 - c. (Optional) For **Alt text**, add the alt text that you want the visual to have.

6. Expand the **Data point limit** section. For **Number of data points to show**, enter the number of data points that you want the visual to show. Highcharts visuals can show up to 10,000 data points.
7. Expand the **Chart code** section.
8. Enter a JSON schema into the **Chart code** JSON editor. The editor provides contextual assistance and real time validation to ensure that your input JSON is configured properly. Any errors that Quick identifies can be viewed in the **Errors** dropdown. The example below shows a JSON schema that creates a lipstick chart that shows current year sales by industry.

```
{
  "xAxis": {
    "categories": ["getColumn", 0]
  },
  "yAxis": {
    "min": 0,
    "title": {
      "text": "Amount ($)"
    }
  },
  "tooltip": {
    "headerFormat": "<span style='font-size:10px'>{point.key}</span><table>",
    "pointFormat": "<tr><td style='color:{series.color};padding:0'>{series.name}:</td><td style='padding:0'><b>${point.y:.0f}</b></td></tr>",
    "footerFormat": "</table>",
    "shared": true,
    "useHTML": true
  },
  "plotOptions": {
    "column": {
      "borderWidth": 0,
      "grouping": false,
      "shadow": false
    }
  },
  "series": [
    {
      "type": "column",
      "name": "Current Year Sales",
      "color": "rgba(124,181,236,1)",
      "data": ["getColumn", 1],
      "pointPadding": 0.3,
      "pointPlacement": 0.0
    }
  ]
}
```

```
}  
]  
}
```

9. Choose **APPLY CODE**. Quick converts the JSON schema into a visual that appears in the analysis. To make changes to the rendered visual, update the appropriate properties in the JSON schema and choose **APPLY CODE**.
10. (Optional) Open the **Reference** dropdown to access links to helpful Highcharts reference material.

When you are happy with the rendered visual, close the properties pane. For more information about Quick Sight specific expressions that can be used to configure a Highcharts visual, see [Amazon Quick JSON expression language for Highcharts visuals](#).

Interactive Highchart features

Highchart visualizations in Amazon Quick Sight support custom actions, highlighting, and custom field color consistencies, allowing you to create interactive and visually cohesive charts that integrate seamlessly with other Quick Sight visuals.

Custom actions

With custom actions, you can define specific behaviors for any data point in your Highchart visualizations. This feature seamlessly integrates with Quick Sight's existing action framework, enabling you to create interactive charts that respond to user clicks. The system currently supports single data point selection, giving you precise control over user interactions. Custom actions can be implemented across various chart types, including line charts, bar charts, and stacked bar charts, among others.

To implement custom actions, you'll need to modify your Highcharts JSON configuration. Add an event block to your series configuration, specifying the click event and the corresponding action. For example:

```
{  
  "series": [{  
    "type": "line",  
    "data": ["getColumn", 1],  
    "name": "value",  
    "events": {
```

```

    "click": [
      "triggerClick", { "rowIndex": "point.index" }
    ]
  }
}]

```

This configuration enables click events on your chart's data points, allowing Quick Sight to handle custom actions based on the selected data.

Cross-visual highlighting

Cross-visual highlighting enhances the interactivity of your dashboards by creating visual connections between different charts. When a user selects elements in one chart, related elements in other visuals are automatically highlighted, while unrelated elements are dimmed. This feature helps users quickly identify relationships and patterns across multiple visualizations, improving data comprehension and analysis.

To enable cross-visual highlighting and maintain field color consistency, use the `quicksight` clause in your Highcharts JSON configuration. This clause acts as a bridge between Highcharts rendering and Quick's visual interaction system. Here's an example of how to set it up:

```

{
  "quicksight": {
    "pointRender": ["updatePointAttributes", {
      "opacity": ["case",
        ["dataMarkMatch", ["getColumnName", 0], "series.name"],
        1, // Full opacity for matching elements
        0.1 // Dim non-matching elements
      ],
      "color": ["getColumnColorOverrides", ["getColumnName", 0], "series.name"]
    }]
  }
}

```

This configuration uses Quick Sight's JSON expression language to dynamically modify visual properties like opacity and color based on user interactions and predefined color schemes.

For more complex scenarios, you can set up highlighting based on multiple conditions. This allows for more nuanced interactivity in your visualizations. The following example highlights elements based on either the quarter or day of the week:

```
{
  "quicksight": {
    "pointRender": ["updatePointAttributes", {
      "opacity": ["case",
        ["||",
          ["dataMarkMatch", "quarter", "series.name"],
          ["dataMarkMatch", "day_of_week", "point.name"]
        ],
        1, // Full opacity for matching elements
        0.1 // Dim non-matching elements
      ],
    }]
  }
}
```

Field-level color consistency

Maintaining visual coherence across your dashboard is crucial for effective data interpretation. The field-level color consistency feature ensures that colors assigned to specific dimensions persist across all visuals in your dashboard. This consistency helps users quickly recognize and track particular data categories across different chart types and views, enhancing the overall user experience and data comprehension.

Amazon Quick JSON expression language for Highcharts visuals

Highcharts visuals accept most [valid JSON values](#), standard arithmetic operators, string operators, and conditional operators. The following JSON values are not supported for Highcharts visuals:

- Functions
- Dates
- Undefined values

Quick authors can use JSON expression language create JSON schemas for a highcharts visual. JSON expression language is used to bind JSON to APIs or datasets to allow dynamic population and modification of JSON structures. Developers can also use JSON expression language to inflate and transform JSON data with concise and intuitive expressions.

With JSON expression language, expressions are represented as arrays, where the first element specifies the operation and subsequent elements are the arguments. For example, ["unique",

[1, 2, 2]] applies the unique operation to the array [1, 2, 2], resulting in [1, 2]. This array-based syntax allows for flexible expressions, that allow complex transformations on JSON data.

JSON expression language supports *nested expressions*. Nested expressions are expressions that contain other expressions as arguments. For example ["split", ["toUpper", "hello world"], " "] first converts the string hello world into an uppercase, then splits it into array of words, resulting in ["HELLO", "WORLD"].

Use the following sections to learn more about JSON expression language for Highcharts visuals in Amazon Quick.

Topics

- [Arithmetics](#)
- [Array operations](#)
- [Amazon Quick expressions](#)

Arithmetics

The following table shows arithmetic expressions that can be used with JSON expression language.

Operation	Expression	Input	Output
Addition	["+", operand1, operand2]	{ sum: ["+", 2, 4] }	{ sum: 6 }
Subtraction	["-", operand1, operand2]	{ difference: ["-", 10, 3] }	{ difference: 7 }
Multiplication	["*", operand1, operand2]	{ product: ["*", 5, 6] }	{ product: 30 }

Operation	Expression	Input	Output
Division	<code>["/", operand1, operand2]</code>	<code>{ quotient: ["/", 20, 4] }</code>	<code>{ quotient: 5 }</code>
Modulo	<code>["%", operand1, operand2]</code>	<code>{ remainder: ["%", 15, 4] }</code>	<code>{ remainder: 3 }</code>
Exponentiation	<code>["**", base, exponent]</code>	<code>{ power: ["**", 2, 3] }</code>	<code>{ power: 8 }</code>
Absolute Value	<code>["abs", operand]</code>	<code>{ absolute: ["abs", -5] }</code>	<code>{ absolute: 5 }</code>
Square Root	<code>["sqrt", operand]</code>	<code>{ sqroot: ["sqrt", 16] }</code>	<code>{ sqroot: 4 }</code>
Logarithm (base 10)	<code>["log10", operand]</code>	<code>{ log: ["log10", 100] }</code>	<code>{ log: 2 }</code>
Natural Logarithm	<code>["ln", operand]</code>	<code>{ ln: ["ln", Math.E] }</code>	<code>{ ln: 1 }</code>
Round	<code>["round", operand]</code>	<code>{ rounded: ["round", 3.7] }</code>	<code>{ rounded: 4 }</code>
Floor	<code>["floor", operand]</code>	<code>{ floor: ["floor", 3.7] }</code>	<code>{ floor: 3 }</code>
Ceiling	<code>["ceil", operand]</code>	<code>{ ceiling: ["ceil", 3.2] }</code>	<code>{ ceiling: 4 }</code>

Operation	Expression	Input	Output
Sine	<code>["sin", operand]</code>	<code>{ sine: ["sin", 0] }</code>	<code>{ sine: 0 }</code>
Cosine	<code>["cos", operand]</code>	<code>{ cosine: ["cos", 0] }</code>	<code>{ cosine: 1 }</code>
Tangent	<code>["tan", operand]</code>	<code>{ tangent: ["tan", Math.PI] }</code>	<code>{ tangent: 0 }</code>

Array operations

JSON expression language allows generic array manipulation for the following functions:

- **map** – Applies a mapping function to each element of an array and returns a new array with the transformed values.

For example, `["map", [1, 2, 3], ["*", ["item"], 2]]` maps each element of the array `[1, 2, 3]` by multiplying it by 2.

- **filter** – Filters an array based on a given condition and returns a new array containing only the elements that satisfy the condition

For example, `["filter", [1, 2, 3, 4, 5], ["==", ["%", ["item"], 2], 0]]` filters the array `[1, 2, 3, 4, 5]` to include only the even numbers.

- **reduce** – Reduces an array to a single value by applying a reducer function to each element and accumulating the result.

For example, `["reduce", [1, 2, 3, 4, 5], ["+", ["acc"], ["item"]], 0]` reduces the array `[1, 2, 3, 4, 5]` to the sum of its elements.

- **get** – Retrieves a value from an object or an array by specifying a key or index.

For example, `["get", ["item"], "name"]` retrieves the value of the "name" property from the current item.

- **unique** – Given an array returns only unique items inside this array.

For example, `["unique", [1, 2, 2]]` returns `[1, 2]`.

Amazon Quick expressions

Amazon Quick offers additional expressions to enhance the functionality of Highcharts visuals. Use the following sections to learn more about common Quick expressions for highcharts visuals. For more information about JSON expression language in Amazon Quick, see the [Highcharts Visual QuickStart Guide](#) in [DemoCentral](#).

Topics

- [getColumn](#)
- [formatValue](#)

getColumn

Use the getColumn expressions to return values from specified column indices. For example, the following table shows a list of products alongside their category, and price.

Product name	Category	Price
Product A	Technology	100
Product B	Retail	50
Product C	Retail	75

The following getColumn query generates an array that shows all product names alongside their price.

```
{
  product name: ["getColumn", 0],
  price: ["getColumn", 2]
}
```

The following JSON is returned:

```
{
```

```
product name: ["Product A", "Product B", "Product C"],  
price: [100, 50, 75]  
}
```

You can also pass multiple columns at once to generate an array of arrays, shown in the following example.

Input

```
{  
  values: ["getColumn", 0, 2]  
}
```

Output

```
{  
  values: [["Product A", 100], ["Product B", 50], ["Product C", 75]]  
}
```

Similar to `getColumn`, the following expressions can be used to return column values from field wells or themes:

- `getColumnFromGroupBy` returns columns from the group by field. The second argument is the index of the column to return. For example, `["getColumnFromGroupBy", 0]` returns values of the first field as an array. You can pass multiple indices to get an array of arrays where each element corresponds to the field in the group by field well.
- `getColumnFromValue` returns columns from the value field well. You can pass multiple indices to get an array of arrays where each element corresponds to the field in the values field well.
- `getColorTheme` returns the current color palette of a Quick theme, shown in the following example.

```
{  
  "color": ["getColorTheme"]  
}
```

```
{  
  "color": ["getPaletteColor", "secondaryBackground"]  
}
```

Example

Visuals

+ ADD

BUILD

CHANGE VISUAL TYPE

Table

GROUP BY 2

day_of_month

day_of_week

Add a dimension

VALUE

cancelled (Sum)

Add a measure

Sheet 1

+

Sum of Cancelled by Day

day_of_mo...	day_of_week	cancelled
1	1	1,674
2	1	7,425
3	1	4,371
4	1	2,657
5	1	1,911
6	1	4,941
7	1	2,151
8	1	3,004
9	1	4,913
10	1	4,876
11	1	2,404
12	1	3,549
13	1	2,845
14	1	2,831
15	1	2,601
16	1	3,198

Page size

500

<< < Page 1 > >>

getColumn can access any column from the table:

- ["getColumn", 0] - returns array [1, 2, 3, 4, 5, ...]
- ["getColumn", 1] - returns array [1, 1, 1, 1, 1, ...]

- ["getColumn", 2] - returns array [1674, 7425, 4371, ...]

getColumnFromGroupBy works similarly, but its index is limited to the columns in the group by field well:

- ["getColumnFromGroupBy", 0] - returns array [1, 2, 3, 4, 5, ...]
- ["getColumnFromGroupBy", 1] - returns array [1, 1, 1, 1, 1, ...]
- ["getColumnFromGroupBy", 2] - does not work, since there are only two columns in the group by field well

getColumnFromValue works similarly, but its index is limited to the columns in the value field well:

- ["getColumnFromValue", 0] - returns array [1, 2, 3, 4, 5, ...]
- ["getColumnFromValue", 1] - does not work, since there is only one column in the value field well
- ["getColumnFromValue", 2] - does not work, since there is only one column in the value field well

formatValue

Use the formatValue expression to apply Quick formatting to your values. For example, the following expression formats the x-axis label with the format value that is specified in the first field of Quick field wells.

```
"xAxis": {
  "categories": ["getColumn", 0],
  "labels": {
    "formatter": ["formatValue", "value", 0]
  }
}
```

Using histograms

Use a histogram chart in Amazon Quick to display the distribution of continuous numerical values in your data. Amazon Quick uses un-normalized histograms, which use an absolute count of the data points or events in each bin.

To create a histogram, you use one measure. A new histogram initially displays ten *bins* (also called *buckets*) across the X-axis. These appear as bars on the chart. You can customize the bins to suit your dataset. The Y-axis displays the absolute count of the values in each bin.

Make sure that you adjust the format settings so that you have a clearly identifiable shape. If your data contains outliers, this becomes clear if you spot one or more values off to the side of the X-axis. For information about how Amazon Quick handles data that falls outside display limits, see [Display limits](#).

Histogram features

To understand the features supported by histograms, use the following table.

Feature	Supported?	Comments	For more information
Changing the legend display	No		Legends on visual types in Quick
Changing the title display	Yes		Titles and subtitles on visual types in Quick
Changing the axis range	No	However, you can change the bin count or the bin interval width (range of distribution).	
Showing or hiding axis lines, grid lines, axis labels, and axis sort icons	Yes		Axes and grid lines on visual types in Quick
Changing the visual colors	Yes		Colors in visual types in Quick
Focusing on or excluding elements	No		
Sorting	No		

Feature	Supported?	Comments	For more information
Performing field aggregation	No	Histograms use only the count aggregation.	
Adding drill-downs	No		

Creating a histogram

Use the following procedure to create a histogram.

To create a histogram

1. On the analysis page, choose **Visualize** on the tool bar.
2. Choose **Add** on the application bar, and then choose **Add visual**.
3. On the **Visual types** pane, choose the histogram icon.
4. On the **Fields list** pane, choose the field that you want to use in the **Value** field well. A **Count** aggregate is automatically applied to the value.

The resulting histogram shows the following:

- The X-axis displays 10 bins by default, representing the intervals in the measure that you choose. You can customize the bins in the next step.
 - The Y-axis displays the absolute count of individual values in each bin.
5. (Optional) Choose **Format** on the visual control to change the histogram format. You can format the bins either by count or width, not both together. The count setting changes how many bins display. The width setting changes how wide or long of an interval each bin contains.

Formatting a histogram

Use the following procedure to format a histogram.

To format a histogram

1. Choose the histogram chart that you want to work with. It should be the highlighted selection. The visual controls display on the top right of the histogram.
2. Choose the cog icon on the visual control menu to view the **Format visual** options.
3. On the **Properties** pane, set the following options to control the display of the histogram:

- **Histogram** settings. Chose *one* of the following settings:
 - Bin count (option 1): The number of bins that display on the X-axis.
 - Bin width (option 1): The width (or length) of each interval. This setting controls the number of items or events to include in each bin. For example, if your data is in minutes, you can set this to 10 to show 10-minute intervals.
- With the following settings, you can explore the best way to format the histogram for your dataset. For example, in some cases, you might have a tall peak in one bin, while most of the other bins look sparse. This isn't a useful view. You can use the following settings individually or together:
 - Change the **Number of data points displayed** in the **X-axis** settings.

Amazon Quick displays up to 100 bins (buckets) by default. If you want to display more (up to 1,000), change the X-axis setting for **Number of data points displayed**.

- Enable **Logarithmic scale** in the **Y-axis** settings.

Sometimes your data doesn't fit the shape that you want and this can provide misleading results. For example, if the shape is skewed so far to the right that you can't read it properly, you can apply a log scale to it. Doing this doesn't normalize your data; however, it does reduce the skew.

- Display **Data labels**.

You can enable the display of data labels to see the absolute counts in the chart. Even if you don't want to display these in most cases, you can enable them while you're developing an analysis. The labels can help you decide on formatting and filtering options because they reveal counts in bins that are too small to stand out.

To see all the data labels, even if they overlap, enable **Allow labels to overlap**.

4. (Optional) Change other visual settings. For more information, see [Formatting in Amazon Quick](#).

Understanding histograms

Although histograms look similar to bar charts, they are very different. In fact, the only similarity is their appearance because they use bars. On a histogram, each bar is called a *bin* or a *bucket*.

Each bin contains a range of values called an *interval*. When you pause on one of the bins, details about the interval appear in a tooltip that shows two numbers enclosed in glyphs. The type of enclosing glyphs indicates if the numbers inside them are part of the interval that's inside the selected bin, as follows:

- A square bracket next to a number means that the number is included.
- A parenthesis next to a number means that the number is excluded.

For example, let's say that the first bar in a histogram displays the following notation.

$[1, 10)$

The square bracket means that the number 1 is included in the first interval. The parenthesis means that the number 10 is excluded.

In the same histogram, a second bar displays the following notation.

$[10, 20)$

In this case, 10 is included in the second interval, and 20 is excluded. The number 10 can't exist in both intervals, so the notation shows us which one includes it.

Note

The pattern used for marking intervals in a histogram comes from standard mathematical notation. The following examples show the possible patterns, using a set of numbers that includes 10, 20, and every number in between.

- $[10, 20]$ – This set is closed. It has hard boundaries on both ends.
- $[10, 21)$ – This set is half open. It has a hard boundary on the left and a soft boundary on the right.
- $(9, 20]$ – This set is half open. It has a soft boundary on the left and a hard boundary on the right.

- (9, 21) – This set is open. It has soft boundaries on both ends.

Because the histogram uses quantitative data (numbers) rather than qualitative data, there's a logical order to the distribution of the data. This is called a *shape*. The shape is often described the qualities the shape possesses, based on the count in each bin. Bins that contain a higher number of values form a *peak*. Bins that contain a lower number of values form a *tail* on the edge of a chart, and a *valley* between peaks. Most histograms fall into one of the following shapes:

- Asymmetrical or *skewed* distributions have values that cluster near the left or the right—the low or high end of the X-axis. The direction of skewness is defined by where the longer tail of the data is, not by where the peak is. It's defined this way because this direction also describes the location of the mean (average). In skewed distributions, the mean and the median are two different numbers. The different types of skewed distribution are as follows:
 - *Negatively skewed* or *left skewed* – A chart that has the mean to the left of the peak. It has a longer tail to the left and a peak to the right, sometimes followed by a shorter tail.
 - *Positively skewed* or *right skewed* – A chart that has the mean to the right of the peak. It has a longer tail to the right and a peak to the left, sometimes preceded by a shorter tail.
- Symmetrical or *normal* distributions have a shape that's mirrored on each side of a center point (for example, a bell curve). In a normal distribution, the mean and the median are the same value. The different types of normal distribution are as follows:
 - Normal distribution, or *unimodal* – A chart that has one central peak representing the most common value. This is commonly called a bell curve, or a Gaussian distribution.
 - Bimodal – A chart that has two peaks representing the most common values.
 - Multimodal – A chart that has three or more peaks representing the most common values.
 - Uniform – A chart that has no peaks or valleys, with a relatively equal distribution of data.

The following table shows how a histogram differs from a bar chart.

Histogram	Bar chart
A histogram displays the distribution of values in one field.	A bar chart compares the values in one field, grouped by dimension.

Histogram	Bar chart
A histogram sorts values into bins that represent a range of values, for example 1–10, 10–20, and so on.	A bar chart plots values that are grouped into categories.
The sum of all bins equals exactly 100% of the values in the filtered data.	A bar chart isn't required to display all of the available data. You can change display settings at the visual level. For example, a bar chart might show only the top 10 categories of data.
Rearranging bars detracts from the meaning of the chart as a whole.	Bars can be in any order without changing the meaning of the chart as a whole.
There are no spaces between the bars, to represent the fact this is continuous data.	There are spaces between the bars, to represent the fact that this is categorical data.
If a line is included in a histogram, it represents the general shape of the data.	If a line is included in a bar chart, it's called a combo chart, and the line represents a different measure than the bars.

Using image components

Use image components to upload static images from your desktop to a Quick analysis. Each Quick analysis sheet supports up to 10 image components. Image components are not included in the 50 visual per sheet limit. The file size of an image component can't exceed 1MB.

The following file formats are supported for image components:

- .bmp
- .jpg/.jpeg
- .png
- .tiff
- .webp

Use the following procedure to add an image component to a Quick analysis:

To add an image component to a Quick analysis

1. Open the [Quick console](#).
2. Open the Quick analysis that you want to add an image to.
3. Choose the **Add image** button in the toolbar at the top of the analysis.
4. The file selection of your desktop opens. Choose the file that you want to upload, and then choose **Open**. The file size of the image component can't exceed 1MB.
5. The image is uploaded to Quick and appears in the analysis.
6. (Optional) To add alt text or update the image scaling options, choose the **Properties** icon at the top right of the image to open the **Properties** pane.
7. (Optional) To add a [custom tooltip](#) to the image, open the **Properties** pane, choose **Interactions**, and then choose **Add action**. Filter actions are not supported for image components. You can also use the **Interactions** section to add custom navigation and URL actions to the image component.
8. (Optional) To duplicate or replace the image, choose the **More options** ellipsis (three dots) icon at the top right of the image, and then choose the action that you want to perform.

Using KPIs

Use a key performance indicator (KPI) to visualize a comparison between a key value and its target value.

A KPI displays a value comparison, the two values being compared, and a visual that provides context to the data that's displayed. You can choose from a set of predesigned layouts to suit your business needs. The following image shows an example of a KPI visual that uses a sparkline.

1. Choose **Add (+)** drop down in the **Visuals** pane.
2. Choose the KPI icon from Visual types menu.

KPI features

To understand the features supported by the KPI visual type in Amazon Quick, use the following table.

Feature	Supported?	Comments	For more information
Changing the title display	Yes		Titles and subtitles on visual types in Quick
Removing the title	Yes	You can choose not to display a title.	
Changing comparison method	Yes	By default, Amazon Quick automatically chooses a method. The settings are auto, difference, percent, and difference as percent.	
Changing the primary value displayed	Yes	You can choose comparison (default) or actual.	
Displaying or removing the progress bar	Yes	You can format the visual to either display (default) or not display a progress bar.	

For more information on KPI formatting options, see [KPI options](#).

Creating a KPI

Use the following procedure to create a KPI.

To create a KPI

1. Create a new analysis for your dataset.
2. In the **Visual types** pane, choose the KPI icon.
3. From the **Fields list** pane, drag the fields that you want to use to the appropriate field wells. You must use measure fields as indicated by the target field well. If you choose to use a dimension field as a measure, the **Count** aggregate function is automatically applied to it to create a numeric value.

To create a KPI, drag a measure to the **Value** field well. To compare that value to a target value, drag a different measure to the **Target value** field well.

4. (Optional) Choose formatting options by selecting the on-visual menu at the upper-right corner of the visual, then choosing **Format visual**.

Changing a KPI's layout

Use the following procedure to change the layout for a KPI.


To change the layout of a KPI

1. Navigate to the KPI visual that you want to change and choose **KPI layouts**.
2. In the **KPI Layouts** pane, choose the KPI layout that you want to use.

Using layer maps

Use layer maps to visualize data with custom geographic boundaries, such as congressional districts, sales territories, or user-defined regions. With layer maps, Quick authors upload GeoJSON files to Amazon Quick that shape layers over a base map and join with Quick data to visualize associated metrics and dimensions. Shape layers can be styled by color, border, and opacity. Quick authors can also add interactivity to layer maps through tooltips and custom actions.

The following image shows a layer map visual in Amazon Quick.



Creating a shape layer with layered maps

Use the procedure below to create a shape layer with layer map visuals in Amazon Quick.

1. Open the [Quick console](#).
2. Open the Quick analysis that you want to add a layer map to.
3. Choose **Add** on the application bar, and then choose **Add visual**.
4. On the **Visual types** pane, choose one of the layer map icon.
5. An empty map visual appears in the analysis and prompts you to continue configuring the layer. Choose **CONFIGURE LAYER** to continue configuring the layer map.
6. The **Layer properties** pane opens to the right. Navigate to the **Shape file** section, and then choose **UPLOAD SHAPE FILE**.

7. Choose the GeoJSON file that you want to visualize. The file must be in .geojson format and must not exceed 100 MB.
8. Navigate to the **Data** section.
9. For **Shape file key field**, choose the field that you want the shape to visualize.
10. (Optional) For **Dataset key field**, choose the dataset field that you want the shape to visualize. To assign color to the shapes, add a color field. If the color field is a measure, the shape uses gradient coloring. If the color field is a dimension, the shape uses categorical coloring. If a color field is not assigned to the shape, use the fill color option in the **Styling** section of the **Layer properties** pane to set a common color for all shapes.
11. (Optional) To change the layer name, navigate to the **Layer options** section and enter a name in the **Layer name** input.
12. (Optional) To change the fill or border colors, navigate to the **Styling** section and choose the color switch next to the object that you want to change. To adjust the opacity of the color, enter a percentage amount in the input located next to the eye icon. If you do not assign a color field to the **Dataset key field**, the fill color can be used to set a common color for all shapes.

Using line charts

Use line charts to compare changes in measure values over period of time, for the following scenarios:

- One measure over a period of time, for example gross sales by month.
- Multiple measures over a period of time, for example gross sales and net sales by month.
- One measure for a dimension over a period of time, for example number of flight delays per day by airline.

Line charts show the individual values of a set of measures or dimensions against the range displayed by the Y axis. Area line charts differ from regular line charts in that each value is represented by a colored area of the chart instead of just a line, to make it easier to evaluate item values relative to each other.

Because a stacked area line chart works differently than other line charts, simplify it if you can. Then the audience won't try to interpret the numbers. Instead, they can focus on the relationships of each set of values to the whole. One way to simplify is to remove the numbers down the left

side of the screen by reducing the step size for the axis. To do this, choose the **Options** icon from the on-visual menu. In **Format Options** under **Y-axis**, enter **2** as the **Step size**.

Each line on the chart represents a measure value over a period of time. You can interactively view the values on the chart. Hover over any line to see a pop-up legend that shows the values for each line on the **X axis**. If you hover over a data point, you can see the **Value** for that specific point on the **X axis**.

Use line charts to compare changes in values for one or more measures or dimensions over a period of time.

In regular line charts, each value is represented by a line, and in area line charts each value is represented by a colored area of the chart.

Use stacked area line charts to compare changes in values for one or more groups of measures or dimensions over a period of time. Stacked area line charts show the total value for each group on the x-axis. They use color segments to show the values of each measure or dimension in the group.

Line charts show up to 10,000 data points on the x-axis when no color field is selected. When color is populated, line charts show up to 400 data points on the x-axis and up to 25 data points for color. For more information about data that falls outside the display limit for this visual type, see [Display limits](#).

Line chart features

To understand the features supported by line charts, use the following table.

Feature	Supported?	Comments	For more information
Changing the legend display	Yes		Legends on visual types in Quick
Changing the title display	Yes		Titles and subtitles on visual types in Quick
Changing the axis range	Yes	You can set the range for the Y axis.	Range and scale on visual types in Quick
Showing or hiding axis lines,	Yes		Axes and grid lines on visual types in Quick

Feature	Supported?	Comments	For more information
grid lines, axis labels, and axis sort icons			
Adding a second Y-axis	Yes		Creating a dual-axis line chart
Changing the visual colors	Yes		Colors in visual types in Quick
Focusing on or excluding elements	Yes, with exceptions	You can focus on or exclude any line on the chart, except in the following cases: - You create a multi-dimension line chart and use a date field as the dimension for the line color. - You create a measure or multi-measure line chart and use a date field as the dimension for the X axis. In these cases, you can only focus on a line, not exclude it.	Focusing on visual elements Excluding visual elements
Sorting	Yes, with exceptions	You can sort data for numeric measures in the X axis and Value field wells. Other data is automatically sorted in ascending order.	Sorting visual data in Amazon Quick

Feature	Supported?	Comments	For more information
Performing field aggregation	Yes	You must apply aggregation to the field that you choose for the value, and can't apply aggregation to the fields you choose for the X axis and color.	Changing field aggregation
Adding drill-downs	Yes	You can add drill-down levels to the X axis and Color field wells.	Adding drill-downs to visual data in Quick Sight

Creating a line chart

Use the following procedure to create a line chart.


To create a line chart

1. On the analysis page, choose **Visualize** on the tool bar.
2. Choose **Add** on the application bar, and then choose **Add visual**.
3. On the **Visual types** pane, choose one of the line chart icons.
4. From the **Fields list** pane, drag the fields that you want to use to the appropriate field wells. Typically, you want to use dimension or measure fields as indicated by the target field well. If you choose to use a dimension field as a measure, the **Count** aggregate function is automatically applied to it to create a numeric value.
 - To create a single-measure line chart, drag a dimension to the **X axis** field well and one measure to the **Value** field well.
 - To create a multi-measure line chart, drag a dimension to the **X axis** field well and two or more measures to the **Value** field well. Leave the **Color** field well empty.
 - To create a multi-dimension line chart, drag a dimension to the **X axis** field well, one measure to the **Value** field well, and one dimension to the **Color** field well.
5. (Optional) Add drill-down layers by dragging one or more additional fields to the **X axis** or **Color** field wells. For more information about adding drill-downs, see [Adding drill-downs to visual data in Quick Sight](#).

Creating a dual-axis line chart


If you have two or more metrics that you want to display in the same line chart, you can create a dual-axis line chart.

A *dual-axis chart* is a chart with two Y-axes (one axis at the left of the chart, and one axis at the right of the chart). For example, let's say you create a line chart. It shows the number of visitors who signed up for a mailing list and for a free service over a period of time. If the scale between those two measures varies widely over time, your chart might look something like the following line chart. Because the scale between measures varies so greatly, the measure with the smaller scale appears nearly flat at zero.



If you want to show these measures in the same chart, you can create a dual-axis line chart. The following is an example of the same line chart with two Y-axes.

Sum of Free Sign Up and Sum of Mailing List Adds by Date




To create a dual-axis line chart

1. In your analysis, create a line chart. For more information about creating line charts, see [Creating a line chart](#).
2. In the **Value field well**, choose a field drop-down menu, choose **Show on: Left Y-axis**, and then choose **Right Y-axis**.

Or you can create a dual-axis line chart using the **Properties** pane:

- a. On the menu in the upper-right corner of the line chart, choose the **Format visual** icon.
- b. In the **Properties** pane that opens, choose **Data series**.
- c. In the **Data series** section, choose the **Show on right axis** icon for the value that you want to place on a separate axis. Use the search bar to quickly find a value if you need to.



The icon updates to indicate that the value is being shown on the right axis. The chart updates with two axes.

The **Properties** pane updates with the following options:

- To synchronize the Y-axes for both lines back into a single axis, choose **Single Y-axis** at the top of the **Properties** pane.
- To format the axis at the left of the chart, choose **Left Y-axis**.
- To format the axis at the right of the chart, choose **Right Y-axis**.

For more information about formatting axis lines, see [Axes and grid lines](#). For more information about adjusting the range and scale of an axis, see [Range and scale](#).

Creating maps and geospatial charts

You can create two types of maps in Quick: point maps and filled maps. *Point maps* show the difference between data values for each location by size. *Filled maps* show the difference between data values for each location by varying shades of color.

Important

Geospatial charts in Quick currently aren't supported in some AWS Regions, including in China.

For help with geospatial issues, see [Geospatial troubleshooting](#).

Before you get started creating maps, do the following:

- Make sure that your dataset contains location data. *Location data* is data that corresponds to latitudinal and longitudinal values. Location data can include a column for latitude and a column for longitude in your dataset. It can also include a column with city names. Quick can chart latitude and longitude coordinates. It also recognizes geographic components such as country, state or region, county or district, city, and ZIP code or postal code.
- Make sure that your location data fields are marked as geospatial data types.
- Consider creating geographic hierarchies.

For more information about working with geospatial data, including changing field data types and creating geospatial hierarchies, see [Adding geospatial data](#).

To learn more about creating maps in Quick, see the following.

Topics

- [Creating point maps](#)
- [Creating filled maps](#)
- [Interacting with maps](#)

Creating point maps

You can create point maps in Quick to show the difference between data values for each location by size. Each point on this type of map corresponds to a geographic location in your data, such as a country, state or province, or city. The size of the points on the map represents the magnitude of the field in the **Size** field well, in relation to other values in the same field. The color of the points represents the values in the **Color** field well. The field values in the **Color** field well display in the legend, if you choose a field for color.

Use the following procedure to create a point map in Quick.

To create point maps in Quick, make sure that you have the following:

- One geospatial field (such as country, state or region, county or district, city, or ZIP code or postal code). Or you can use one latitude field and one longitude field.

- One numeric field (measure) for size.
- (Optional) A categorical field (dimension) for color.

For information on formatting geospatial maps, see [Map and geospatial chart formatting options](#).

Creating point maps

To create a point map

1. Add a new visual to your analysis. For more information about starting analyses, see [Starting an analysis in Quick Sight](#). For more information about adding visuals to analyses, see [Adding a visual](#).
2. For **Visual type**, choose the **Points on map** icon. It looks like a globe with a point on it.
3. Drag a geographic field from the **Fields list** pane to the **Geospatial** field well, for example **Country**. You can also choose a latitude or longitude field.

A point map appears with a point for each location in your data.

If the field is part of a geographic hierarchy, the hierarchy displays in the field well.

4. Drag a measure from the **Fields list** pane to the **Size** field well.

The points on the map update to show the magnitude of values for each location.

5. (Optional) Drag a dimension from the **Fields list** pane to the **Color** field well.

Each point updates to show a point for each categorical value in the dimension.

Creating filled maps

You can create filled maps in Quick to show the difference between data values for each location by varying shades of color.

Use the following procedure to create a filled map in Quick.

To create filled maps in Quick, make sure that you have the following:

- One geospatial field (such as country, state or region, county or district, or ZIP code or postal code).
- (Optional) A numeric field (measure) for color.

Creating filled maps

To create a filled map

1. Add a new visual to your analysis. For more information about starting analyses, see [Starting an analysis in Quick Sight](#). For more information about adding visuals to analyses, see [Adding a visual](#).
2. For **Visual type**, choose the **Filled map** icon.
3. Drag a geographic field from the **Fields list** pane to the **Location** field well, for example **Country**.

A filled map appears with each location in your data filled in by the number of times they appear in your dataset (the count).

If the field is part of a geographic hierarchy, the hierarchy displays in the field well.

4. (Optional) Drag a measure from the **Fields list** pane to the **Color** field well, for example **Sales**.

Each location updates to show the sum of sales.

Interacting with maps

When you view a map visual in an Quick analysis or published dashboard, you can interact with it to explore your data. You can pan, zoom in and out, and autozoom to all the data.

By default, map visuals are always zoomed based on the underlying data. When you pan around in the map or zoom to a different level, the zoom to data icon appears above the zoom in and out icons at bottom right of the map. Using this option, you can quickly zoom back to the underlying data.

To pan in a map visual

- Click anywhere on the map visual and drag your cursor in the direction that you want to pan the map.

To zoom in or out in a map visual

- On the map visual, choose the plus or minus icons at bottom right. Or you can double-click the map to zoom in, and shift-double-click to zoom out.

To zoom back to all the data

- On the map visual, choose the zoom to data icon. This icon appears when you pan or zoom in on a map.

Using small multiples

Use this feature when you need to set multiple comparative visuals in a row. When you activate the *small multiples* feature, Amazon Quick creates a container or shelf of small visuals, presented side-by-side. Each copy of the visual contains a one view of the data. Using small multiples is a way to get a holistic view of your business, in an efficient and interactive way.

Small multiples aren't listed in the palette visualization icons. Instead, the option to create small multiples appears as a field well, in the visuals that support it.

To add small visuals to your analysis

1. On a line, bar, or pie charts, add a field to the **Small multiples** field well.
2. To see your small multiples, you need to enlarge the container that holds them, so you can see all of them at once.
3. To format the set of small multiples, choose Format visual (the pencil icon) from the menu on the visual. You can adjust the following settings:

- **Layout**

- **Visible rows**
- **Visible columns**
- **Number of panels**
- Panel title options (toggle)
 - Font size and color
 - Font weight
 - Text alignment
- **Panel order options (toggle)**

Line thickness, style, and color

- **Panel gutter (toggle)**

Spacing

- **Panel background** (toggle)

Background color

Using pie charts

Use pie charts to compare values for items in a dimension. The best use for this type of chart is to show a percentage of a total amount.

Each wedge in a pie chart represents one item in the dimension. Wedge size represents the proportion of the value for the selected measure that the item represents compared to the whole for the dimension. Pie charts are best when precision isn't important and there are few items in the dimension.

To create a donut chart, use one dimension in the **Group/Color** field well. With only one field, the chart displays the division of values by row count. To display the division of dimension values by a metric value, you can add a metric field to the **Value** field well.

Pie charts show up to 20 data points for group or color. For more information about how Amazon Quick handles data that falls outside display limits, see [Display limits](#).

Pie chart features

To understand the features supported by pie charts, use the following table.

Feature	Supported?	Comments	For more information
Changing the legend display	Yes		Legends on visual types in Quick
Changing the title display	Yes		Titles and subtitles on visual types in Quick
Changing the axis range	Not applicable		Range and scale on visual types in Quick
Showing or hiding axis labels.	Yes		Axes and grid lines on visual types in Quick

Feature	Supported?	Comments	For more information
Changing the visual colors	Yes		Colors in visual types in Quick
Focusing on or excluding elements	Yes, with exceptions	You can focus on or exclude a wedge in a pie chart, except when you are using a date field as a dimension . In that case, you can only focus on a wedge, not exclude it.	Focusing on visual elements Excluding visual elements
Sorting	Yes	You can sort on the field that you choose for the value or the group or color.	Sorting visual data in Amazon Quick
Performing field aggregation	Yes	You must apply aggregation to the field that you choose for the value, and can't apply aggregation to the field that you choose for group or color.	Changing field aggregation
Adding drill-downs	Yes	You can add drill-down levels to the Group/Color field well.	Adding drill-downs to visual data in Quick Sight

Creating a pie chart

Use the following procedure to create a pie chart.

To create a pie chart

1. On the analysis page, choose **Visualize** on the tool bar.
2. Choose **Add** on the application bar, and then choose **Add visual**.
3. On the **Visual types** pane, choose the pie chart icon.

4. From the **Fields list** pane, drag the fields that you want to use to the appropriate field wells. Typically, you want to use dimension or measure fields as indicated by the target field well. If you choose to use a dimension field as a measure, the **Count** aggregate function is automatically applied to it to create a numeric value.

To create a pie chart, drag a dimension to the **Group/Color** field well. Optionally, drag a measure to the **Value** field well.

5. (Optional) Add drill-down layers by dragging one or more additional fields to the **Group/Color** field well. For more information about adding drill-downs, see [Adding drill-downs to visual data in Quick Sight](#).

Using pivot tables

Use pivot tables to show measure values for the intersection of two dimensions.

Heat maps and pivot tables display data in a similar tabular fashion. Use a heat map if you want to identify trends and outliers, because the use of color makes these easier to spot. Use a pivot table if you want to analyze data on the visual.

To create a pivot table, choose at least one field of any data type, and choose the pivot table icon. Amazon Quick creates the table and populates the cell values with the count of the column value for the intersecting row value. Typically, you choose a measure and two dimensions measurable by that measure.

Pivot tables support scroll down and right. You can add up to 20 fields as rows and 20 fields as columns. Up to 500,000 records are supported.

Using a pivot table, you can do the following:

- Specify multiple measures to populate the cell values of the table, so that you can see a range of data
- Cluster pivot table columns and rows to show values for subcategories grouped by related dimension
- Sort values in pivot table rows or columns
- Apply statistical functions
- Add totals and subtotals to rows and columns
- Use infinite scroll

- Transpose fields used by rows and columns
- Create custom total aggregations

To easily transpose the fields used by the rows and columns of the pivot table, choose the orientation icon



near the top right of the visual. To see options for showing and hiding totals and subtotals, formatting the visual, or exporting data to a CSV file, choose the Menu items icon at top right.

As with all visual types, you can add and remove fields. You can also change the field associated with a visual element, change field aggregation, and change date field granularity. In addition, you can focus on or exclude rows or columns. For more information about how to make these changes to a pivot table, see [Changing fields used by a visual in Amazon Quick](#).

For information on formatting pivot tables, see [Formatting in Amazon Quick](#).

For information on custom total aggregations for pivot tables, see [Custom total values](#).

Topics

- [Pivot table features](#)
- [Creating a pivot table](#)
- [Orienting pivot table values](#)
- [Expanding and collapsing pivot table clusters](#)
- [Showing and hiding pivot table columns in Quick](#)
- [Sorting pivot tables in Quick](#)
- [Using table calculations in pivot tables](#)
- [Pivot table limitations](#)
- [Pivot table best practices](#)

Pivot table features

Pivot tables don't display a legend.

To understand the features supported by pivot tables, use the following table.

Feature	Supported?	Comments	For more information
Changing the legend display	No		Legends on visual types in Quick
Changing the title display	Yes		Titles and subtitles on visual types in Quick
Changing the axis range	Not applicable		Range and scale on visual types in Quick
Changing the visual colors	No		Colors in visual types in Quick
Focusing on or excluding elements	Yes, with exceptions	You can focus on or exclude any column or row, except when you are using a date field as one of the dimension s. In that case, you can only focus on the column or row that uses the date dimension , not exclude it.	Focusing on visual elements Excluding visual elements
Sorting	Yes	You can sort fields in the Rows or Columns field wells alphabetically or by a metric in ascending or descending order.	Sorting visual data in Amazon Quick Sorting pivot tables in Quick
Performing field aggregation	Yes	You must apply aggregation to the field or fields you choose for the value. You can't apply aggregation to the fields that you choose for the rows or columns. If you choose to create a multi-measure pivot table,	Changing field aggregation

Feature	Supported?	Comments	For more information
		you can apply different types of aggregation to the different measures. For example, you can show the sum of the sales amount and the maximum discount amount.	
Adding drill-downs	No		Adding drill-downs to visual data in Quick Sight
Showing and hiding totals and subtotals	Yes	You can show or hide totals and subtotals for rows and columns. Metrics automatically roll up to show subtotals when you collapse a row or column. If you use a table calculation, use aggregates to display roll-ups.	
Exporting or copying data	Yes	You can export all of the data to a CSV file. You can select and copy the content of the cells.	Exporting data from visuals
Conditional formatting	Yes	You can add conditional formatting for values, subtotals and totals.	Conditional formatting on visual types in Quick

Creating a pivot table

Use the following procedure to create a pivot table.

To create a pivot table

1. On the analysis page, choose the **Visualize** icon on the tool bar.
2. On the **Visuals** pane, choose **+ Add**, and then choose the pivot table icon.
3. From the **Fields list** pane, choose the fields that you want to include. Amazon Quick automatically places these into the field wells.

To change the placement of a field, drag it to the appropriate field wells. Typically, you use dimension or measure fields as indicated by the target field well. If you choose to use a dimension field as a measure, the **Count** aggregate function is automatically applied to it to create a numeric value.

- To create a single-measure pivot table, drag a dimension to the **Rows** field well, a dimension to the **Columns** field well, and a measure to the **Values** field well.
- To create a multi-measure pivot table, drag a dimension to the **Rows** field well, a dimension to the **Columns** field well, and two or more measures to the **Values** field well.
- To create a clustered pivot table, drag one or more dimensions to the **Rows** field well, one or more dimensions to the **Columns** field well, and a measure to the **Values** field well.

You can also select multiple fields for all of the pivot table field wells if you want to. Doing this combines the multi-measure and clustered pivot table approaches.

Note

To view roll-ups for calculated fields, make sure that you are using aggregates. For example, a calculated field with `field-1 / field-2` doesn't display a summary when rolled up. However, `sum(field-1) / sum(field-2)` does display a roll-up summary.

Choosing a layout

When you create a pivot table in Amazon Quick, you can further customize the way your data is presented with Tabular and Hierarchy layout options. For pivot tables that use a tabular layout, each row field is displayed in its own column. For pivot tables that use a hierarchy layout, all row fields are displayed in a single column. Indentation is used to differentiate row headers of different fields. To change the layout of a pivot table visual, open the **Format visual** menu of the pivot

table that you want to change and choose the layout option that you want from the **Pivot options** section.

Depending on the layout that you choose for your pivot table visual, different formatting options are available. For more information about formatting differences between tabular and hierarchy pivot tables, see [Table and pivot table formatting options in Quick](#).

Orienting pivot table values

You can choose to display a pivot table in a columnar or row-based format. Columnar is the default. When you change to a row-based format, a column with the value name is added to the right of the row header column.

To change a pivot table format

1. On the analysis page, choose the pivot table visual that you want to edit.
2. Expand the **Field wells** pane by choosing the field wells at the top of the visual.
3. On the **Values** field well, choose one of the following options:
 - Choose **Column** for a columnar format.
 - Choose **Row** for a row format.

Note

If you use only one metric, you can eliminate the repeated header by formatting the visual and styling it with the **Hide single metric** option.

Expanding and collapsing pivot table clusters

If you are using grouped columns or rows in a pivot table, you can expand or collapse a group to show or hide its data in the visual.

To expand or collapse a pivot table group

1. On the analysis page, choose the pivot table visual that you want to edit.
2. Choose one of the following:
 - To collapse a group, choose the collapse icon near the name of the field.

- To expand a group, choose the expand icon near the name of the field. The collapse icon shows a minus sign. The expand icon shows a plus sign.

In the following screenshot, Customer Region and the Enterprise segment are expanded, and SMB and Startup are collapsed. When a group is collapsed, its data is summarized in the row or column.

Revenue by Service Line, Customer Region, Customer Segment, and Consumption Channel

		Customer Segment > Service Line					
Customer Region	Consumption C...	Billed Amount	Enterprise			SMB	Startup
			Marketing	HR	Billing		
APAC	API	Billed Amount	104,350.34	132,225.58	127,772.08	98,977.33	854,403.2
	Mobile	Billed Amount	191,448.63	247,966.85	238,773.51	231,934.61	1,985,478.41
	Web	Billed Amount	282,733.31	370,020.33	340,658.32	329,184.75	2,854,727.09
EMEA	API	Billed Amount	253,154.34	353,785.82	427,324.86	199,635.93	599,882.97
	Mobile	Billed Amount	445,424.8	663,531.12	787,623.59	474,354.34	1,399,320.6
	Web	Billed Amount	659,433.53	951,639.98	1,136,865.12	667,843.43	2,018,343.87
US	API	Billed Amount	213,715.84	440,195.3	383,297.6	271,066.46	2,193,052.91
	Mobile	Billed Amount	497,424.85	1,029,193.5	899,185.99	656,111.84	5,188,415.35
	Web	Billed Amount	714,712.03	1,466,952.72	1,284,108.35	928,394.53	7,381,337.62

Showing and hiding pivot table columns in Quick

By default, all columns, rows, and their field values appear when you create a pivot table. You can hide columns and rows that you don't want to appear in the pivot table without changing the pivot table values. When you have more than one measure in the pivot table, you can also hide values.

At any time, you can choose to show any hidden fields in the pivot table. When you publish the visual as part of a dashboard, anyone who subscribes to the dashboard can export the pivot table to a comma-separated value (CSV) or Microsoft Excel file. They can choose to export only the visible fields, or all fields. For more information, see [Exporting data from a dashboard to a CSV](#).

To hide a column or row in a pivot table

- In your analysis, select the pivot table visual that you want to work with.
- Choose the three-dot menu in the **Rows**, **Columns** or **Values** field wells, and then choose **Hide**.

To show all hidden fields in a pivot table

1. In your analysis, select the pivot table visual that you want to work with.
2. Choose any field in the **Fields well** and choose **Show all hidden fields**.

Sorting pivot tables in Quick

In Amazon Quick, you can sort values in a pivot table by fields in the **Rows** and **Columns** field wells or quickly by column headers in the pivot table. In pivot tables, you can sort rows and columns independently of each other in alphabetical order, or by a measure.

Note

You can't run Total, Difference, and Percent Difference table calculations when a pivot table is being sorted by a measure. For more information about using table calculations in pivot tables, see [Using table calculations in pivot tables](#).

Understanding sorting in pivot tables

When you have multiple panes in a pivot table, sorting is applied to each pane independently. For example, the Segment column in the pivot table on the left is being sorted in ascending order by Cost. Given that there are multiple panes, the sort starts over for each pane and the rows within each pane (for Segment) are ordered by lowest to highest cost. The table on the right has the same sort applied, but the sort is being applied across the entire table, as shown following.

Pivot Table: Segment sorted by Cost in ascending order

Customer Region	Segment	Cost
APAC	SMB	443,364.19
	Enterprise	1,360,766.43
	Startup	3,852,507.31
EMEA	SMB	961,547.22
	Startup	2,869,902.85
	Enterprise	4,054,584.77
US	SMB	1,191,778.64
	Enterprise	4,445,523.69
	Startup	9,499,922.96

Table: Segment sorted by Cost in ascending order

Customer Region	Segment	Cost
APAC	SMB	443,364.19
EMEA	SMB	961,547.22
US	SMB	1,191,778.64
APAC	Enterprise	1,360,766.43
EMEA	Startup	2,869,902.85
APAC	Startup	3,852,507.31
EMEA	Enterprise	4,054,584.77
US	Enterprise	4,445,523.69
US	Startup	9,499,922.96

When you apply multiple sorts to a pivot table, sorting is applied from the outside dimension to the inside dimension. Consider the following example image of a pivot table. The Customer Region column is sorted by Cost in descending order (as shown in orange). The Channel column is sorted by Revenue Goal in ascending order (as shown in blue).



Customer Region	Channel	Revenue Goal	Cost
US	API	\$3,308,145	\$2,359,267
	Mobile	\$7,445,681	\$5,291,181
	Web	\$9,718,851	\$7,486,778
	Subtotal	\$20,472,677	\$15,137,225
EMEA	API	\$2,023,046	\$1,344,824
	Mobile	\$3,660,777	\$2,933,557
	Web	\$4,895,932	\$3,607,654
	Subtotal	\$10,579,754	\$7,886,035
APAC	API	\$1,392,361	\$1,067,056
	Mobile	\$2,606,978	\$1,876,292
	Web	\$3,400,757	\$2,713,291
	Subtotal	\$7,400,096	\$5,656,638

Sorting pivot tables using row or column headers

Use the following procedure to sort a pivot table using Row or Column headers.

To sort values in a tabular pivot table using table headers

1. In a tabular pivot table chart, choose the header that you want to sort.
2. For **Sort by**, choose a field to sort by and a sort order.

You can sort dimension fields alphabetically a–z or z–a, or you can sort them by a measure in ascending or descending order.

Pivot Table: Segment sorted by Cost in ascending order

Customer Region	Segment	Cost
APAC	Enterprise	1,360,766.43
	SMB	443,364.19
	Startup	3,852,507.31
EMEA	Enterprise	4,054,584.77
	SMB	961,547.22
	Startup	2,869,902.85
US	Enterprise	4,445,523.69
	SMB	1,191,778.64
	Startup	9,499,922.96

Sorting pivot tables using value headers

Use the following procedure to sort a pivot table using value headers.

To sort a pivot table using value headers

1. In a pivot table chart, choose the value header that you want to sort.
2. Choose **Ascending** or **Descending**.

Customer Region	Segment	Revenue Goal	Cost
APAC	Enterprise	1,873,330.74	1,360,766.43
	SMB	523,381.34	443,364.19
	Startup	5,003,383.53	3,852,507.31
EMEA	Enterprise	5,450,312.56	4,054,584.77
	SMB	1,321,251.73	961,547.22
	Startup	3,808,189.88	2,869,902.85
US	Enterprise	6,353,745.66	4,445,523.69
	SMB	1,506,943.36	1,191,778.64
	Startup	12,611,988.15	9,499,922.96

Sorting by value headers in a pivot table also works on subtotals.

Sorting tabular pivot tables using the field wells

Use the following procedure to sort values in a tabular pivot table using the field wells.

To sort values in a tabular pivot table using the field wells

1. On the analysis page, choose the tabular pivot table that you want to sort.
2. Expand the **Field wells**.
3. In the **Rows** or **Columns** field well, choose the field that you want to sort, and then choose how you want to sort the field for **Sort by**.

You can sort dimension fields in the **Rows** or **Columns** field wells alphabetically from a–z or z–a, or you can sort them by a measure in ascending or descending order. You also have the option to collapse all or expand all rows or columns for the field you choose in the field well. You can also remove the field, or to replace it with another field.

- To sort a dimension field alphabetically, hover your cursor over the field in the **Rows** or **Columns** field well, and then choose the a–z or z–a sort icon.

Visuals

+ ADD

CHANGE VISUAL TYPE

Pivot table

ROWS 2

Segment

Industry

Add a dimension

COLUMNS

Add a dimension

VALUES 2

Sales (Sum)

Quantity (Sum)

Sheet 1

Sum of Sales and Sum of Quantity by Segment and Industry

Rows	Sales	Quantity
Enterprise	429,653.15	6,744
Small Business	1,161,401.35	19,521
Mid Market	706,146.37	11,608

Sort by: Segment

Sort order: Ascending

Format: Text

Collapse

Expand

Hide

Remove

Search fields

City

Contact Name

- To sort a dimension field by a measure, hover your cursor over the field in the **Rows** or **Columns** field well. Then choose a measure from the list, and then choose the ascending or descending sort icon.

The screenshot shows the Amazon QuickSight interface. On the left, the 'Visuals' pane is open, showing a 'Pivot table' visual. The 'ROWS' field well contains 'Segment' and 'Industry'. The 'VALUES' field well contains 'Sales (Sum)' and 'Quantity (Sum)'. A context menu is open over the pivot table, showing 'Sort by: Sales', 'Sort order: Ascending', and 'Format: Text'. A red box highlights the 'Sort options' menu, which includes 'Quantity', 'Sales' (selected), and 'Segment'. Below the pivot table, a table of data is visible:

Rows	Sales	Quantity
Enterprise	429,653.15	6,744
...	...	...
...	20,984.49	345

Or, if you want more control over how the sort is applied to the pivot table, customize the sort options.

To create a sort using the sort options

1. On the analysis page, choose the pivot table that you want to sort.
2. Expand **Field wells**.
3. Choose the field that you want to sort in the **Rows** or **Columns** field well, and then choose **Sort options**.
4. In the **Sort options** pane that opens at left, specify the following options:
 - a. For **Sort by**, choose a field from the drop-down list.
 - b. For **Aggregation**, choose an aggregation from the list.
 - c. For **Sort order**, select **Ascending** or **Descending**.
 - d. Choose **Apply**.

Sorting hierarchy pivot tables using the field wells

For tabular pivot tables, each field in the **Rows** field well has a separate title cell. For hierarchy pivot tables, all row fields are displayed in a single column. To sort, collapse, and expand these row fields, select the **Rows** label to open the **Combined row fields** menu and choose the option that you want. Each field in a hierarchy pivot table can be individually sorted from the **Combined row fields** menu.

Region > Ship Mode

	Central			East		
	Regular Air	Express Air	Delivery Truck	Regular Air	Express Air	De
Rows	Sales	Sales	Sales	Sales	Sales	
Expand						
Collapse						
Product Category	Sort by: Product Category	Product Category	17,323.77			
Product Sub-Category	Sort order: Ascending	# Sales				
Product Container	393.38		276,138.02			
Product Name	45,362.14	301.98		14,712.99		
Medium Box	3,331.77	2,382.29		18,724.15	11,350.75	
Office Furnishings						
Jumbo Drum	176.33		298.51			
Large Box	19,088.67	3,470.48		24,270.16	7,495.49	
Medium Box	36,487.98	6,193.95		10,621.23	6,376.6	
Small Box	28,754.9	4,993.22		17,615.19	1,601.23	
Small Pack	20,554.86	1,329.09		12,245.54	2,983.67	
Wrap Bag	12,150.68	2,057.71		10,627.04	479.78	
Tables						
Jumbo Box	95.97		230,167.66			

More advanced formatting options such as **Hide** and **Remove** are available from the field well menus.

Using table calculations in pivot tables

You can use table calculations to apply statistical functions to pivot table cells that contain measures (numeric values). Use the following sections to understand which functions you can use in calculations, and how to apply or remove them.

The data type of the cell value automatically changes to work for your calculation. For example, say that you apply the **Rank** function to a currency data type. The values display as integers rather than currency, because rank isn't measured as currency. Similarly, if you apply the **Percent difference** function instead, the cell values display as percentages.

Topics

- [Adding and deleting pivot table calculations](#)
- [Functions for pivot table calculations](#)
- [Ways to apply pivot table calculations](#)

Adding and deleting pivot table calculations

Use the following procedures to add, modify, and remove table calculation on a pivot table.

Topics

- [Adding a pivot table calculation](#)
- [Changing how a calculation is applied](#)
- [Removing a calculation](#)

Adding a pivot table calculation

Use the following procedure to add a table calculation to a pivot table.

To add a table calculation to a pivot table

1. Expand the **Field wells** pane by choosing the field wells near the bottom of the visual.
2. Choose the field in the **Values** well that you want to apply a table calculation to, choose **Add table calculation**, and then choose the function to apply.

Note

You can't run Total, Difference, and Percent Difference table calculations when a pivot table is being sorted by a measure. To use these table calculations, remove the sort from the pivot table.

Changing how a calculation is applied

Use the following procedure to change the way a table calculation is applied to a pivot table.

To change the way a table calculation is applied to a pivot table

1. Expand the **Field wells** pane by choosing field wells at the top of the visual.
2. Choose the field in the **Values** well that has the table calculation that you want to change, choose **Calculate as**, and then choose the way that you want the calculation applied.

Removing a calculation

Use the following procedure to remove a table calculation from a pivot table.

To remove a table calculation from a pivot table

1. Expand the **Field wells** pane by choosing the field wells near the bottom of the visual.
2. Choose the field in the **Values** well that you want to remove the table calculation from, and then choose **Remove**.

Functions for pivot table calculations

You can use the following functions in pivot table calculations.

Topics

- [Running total](#)
- [Difference](#)
- [Percentage difference](#)
- [Percent of total](#)
- [Rank](#)
- [Percentile](#)

You can apply functions listed to the following data:

Sum of Billed Amount by Date, Customer Region, Consumption Channel, and Service Line							
		Service Line > Consumption Channel					
Customer Region	Date	Billing			HR		
		Web	Mobile	API	Web	Mobile	API
APAC	2014	\$257K	\$178K	\$82K	\$246K	\$163K	\$74K
	2015	\$471K	\$335K	\$149K	\$465K	\$322K	\$146K
	2016	\$819K	\$572K	\$260K	\$805K	\$548K	\$253K
EMEA	2014	\$376K	\$259K	\$122K	\$353K	\$249K	\$114K
	2015	\$589K	\$414K	\$209K	\$576K	\$401K	\$198K
	2016	\$968K	\$671K	\$336K	\$932K	\$651K	\$316K
US	2014	\$709K	\$496K	\$213K	\$729K	\$501K	\$217K
	2015	\$1,253K	\$872K	\$375K	\$1,267K	\$903K	\$375K
	2016	\$2,181K	\$1,538K	\$645K	\$2,192K	\$1,554K	\$654K

Sum of Billed Amount by Date, Customer Region, Consumption Channel, and Service Line							
		Service Line > Consumption Channel					
Customer Region	Date	Billing			HR		
		Web	Mobile	API	Web	Mobile	API
APAC	2014	\$257K	\$178K	\$82K	\$246K	\$163K	\$74K
	2015	\$471K	\$335K	\$149K	\$465K	\$322K	\$146K
	2016	\$819K	\$572K	\$260K	\$805K	\$548K	\$253K
EMEA	2014	\$376K	\$259K	\$122K	\$353K	\$249K	\$114K
	2015	\$589K	\$414K	\$209K	\$576K	\$401K	\$198K
	2016	\$968K	\$671K	\$336K	\$932K	\$651K	\$316K
US	2014	\$709K	\$496K	\$213K	\$729K	\$501K	\$217K
	2015	\$1,253K	\$872K	\$375K	\$1,267K	\$903K	\$375K
	2016	\$2,181K	\$1,538K	\$645K	\$2,192K	\$1,554K	\$654K

Running total

The **Running total** function calculates the sum of a given cell value and the values of all cells prior to it. This sum is calculated as $Cell1=Cell1$, $Cell2=Cell1+Cell2$, $Cell3=Cell1+Cell2+Cell3$, and so on.

Applying the **Running total** function across the table rows, using **Table across** for **Calculate as**, gives you the following results.

Sum of Billed Amount by Date, Customer Region, Consumption Channel, and Service Line							
		Service Line > Consumption Channel					
Customer Region	Date	Billing			HR		
		Web	Mobile	API	Web	Mobile	API
APAC	2014	\$257K	\$435K	\$516K	\$762K	\$925K	\$999K
	2015	\$471K	\$806K	\$954K	\$1,419K	\$1,742K	\$1,888K
	2016	\$819K	\$1,391K	\$1,651K	\$2,456K	\$3,005K	\$3,258K
EMEA	2014	\$376K	\$635K	\$757K	\$1,111K	\$1,360K	\$1,473K
	2015	\$589K	\$1,003K	\$1,212K	\$1,788K	\$2,188K	\$2,386K
	2016	\$968K	\$1,639K	\$1,975K	\$2,907K	\$3,558K	\$3,873K
US	2014	\$709K	\$1,205K	\$1,417K	\$2,146K	\$2,647K	\$2,864K
	2015	\$1,253K	\$2,126K	\$2,500K	\$3,767K	\$4,670K	\$5,044K
	2016	\$2,181K	\$3,720K	\$4,365K	\$6,557K	\$8,111K	\$8,765K

Difference

The **Difference** function calculates the difference between a cell value and value of the cell prior to it. This difference is calculated as $Cell1=Cell1-null$, $Cell2=Cell2-Cell1$, $Cell3=Cell3-Cell2$, and so on. Because $Cell1-null = null$, the Cell1 value is always empty.

Applying the **Difference** function across the table rows, using **Table across** for **Calculate as**, gives you the following results.

Sum of Billed Amount by Date, Customer Region, Consumption Channel, and Service Line							
Service Line > Consumption Channel							
Customer Region	Date	Billing			HR		
		Web	Mobile	API	Web	Mobile	API
APAC	2014		-\$78K	-\$96K	\$164K	-\$83K	-\$90K
	2015		-\$136K	-\$186K	\$317K	-\$143K	-\$176K
	2016		-\$247K	-\$312K	\$545K	-\$257K	-\$295K
EMEA	2014		-\$117K	-\$137K	\$231K	-\$104K	-\$135K
	2015		-\$176K	-\$204K	\$366K	-\$175K	-\$203K
	2016		-\$297K	-\$335K	\$595K	-\$281K	-\$335K
US	2014		-\$213K	-\$284K	\$517K	-\$228K	-\$284K
	2015		-\$381K	-\$498K	\$892K	-\$364K	-\$528K
	2016		-\$643K	-\$893K	\$1,547K	-\$638K	-\$900K

Percentage difference

The **Percentage Difference** function calculates the percent difference between a cell value and the value of the cell prior to it, divided by the value of the cell prior to it. This value is calculated as $Cell1=(Cell1-null)/null$, $Cell2=(Cell2-Cell1)/Cell1$, $Cell3=(Cell3-Cell2)/Cell2$, and so on. Because $(Cell1-null)/null = null$, the Cell1 value is always empty.

Applying the **Percentage Difference** function across the table rows, using **Table across** for **Calculate as**, gives you the following results.

Sum of Billed Amount by Date, Customer Region, Consumption Channel, and Service Line							
Service Line > Consumption Channel							
Customer Region	Date	Billing			HR		
		Web	Mobile	API	Web	Mobile	API
APAC	2014		-30.56%	-54.15%	200.96%	-33.64%	-54.88%
	2015		-28.91%	-55.61%	213.06%	-30.67%	-54.71%
	2016		-30.17%	-54.52%	209.77%	-31.94%	-53.80%
EMEA	2014		-31.03%	-52.93%	189.54%	-29.52%	-54.34%
	2015		-29.79%	-49.37%	174.91%	-30.38%	-50.64%
	2016		-30.68%	-49.86%	177.00%	-30.16%	-51.51%
US	2014		-29.99%	-57.15%	243.08%	-31.27%	-56.68%
	2015		-30.37%	-57.04%	237.94%	-28.71%	-58.50%
	2016		-29.48%	-58.06%	239.84%	-29.10%	-57.92%

Percent of total

The **Percent of Total** function calculates the percentage the given cell represents of the sum of all of the cells included in the calculation. This percentage is calculated as $\text{Cell1} = \text{Cell1} / (\text{sum of all cells})$, $\text{Cell2} = \text{Cell2} / (\text{sum of all cells})$, and so on.

Applying the **Percent of Total** function across the table rows, using **Table across** for **Calculate as**, gives you the following results.

Sum of Billed Amount by Date, Customer Region, Consumption Channel, and Service Line							
Service Line > Consumption Channel							
Customer Region	Date	Billing			HR		
		Web	Mobile	API	Web	Mobile	API
APAC	2014	21.47%	14.90%	6.83%	20.56%	13.64%	6.15%
	2015	20.54%	14.60%	6.48%	20.29%	14.06%	6.37%
	2016	20.40%	14.24%	6.47%	20.06%	13.65%	6.30%
EMEA	2014	20.00%	13.79%	6.49%	18.79%	13.24%	6.04%
	2015	19.07%	13.39%	6.77%	18.63%	12.97%	6.40%
	2016	19.31%	13.38%	6.71%	18.59%	12.98%	6.29%
US	2014	19.61%	13.72%	5.88%	20.17%	13.86%	6.00%
	2015	19.38%	13.49%	5.79%	19.59%	13.96%	5.79%
	2016	19.38%	13.66%	5.73%	19.48%	13.81%	5.81%

Rank

The **Rank** function calculates the rank of the cell value compared to the values of the other cells included in the calculation. Rank always shows the highest value equal to 1 and lowest value equal to the count of cells included in the calculation. If there are two or more cells with equal values, they receive the same rank but are considered to take up their own spots in the ranking. Thus, the next highest value is pushed down in rank by the number of cells at the rank above it, minus one. For example, if you rank the values 5,3,3,4,3,2, their ranks are 1,3,3,2,3,6.

For example, suppose that you have the following data.

Customer Region	Service Line		
	Marketing	HR	Billing
APAC	\$1,357K	\$3,023K	\$3,122K
EMEA	\$2,247K	\$3,788K	\$3,944K
US	\$4,657K	\$8,392K	\$8,282K

Applying the **Rank** function across the table rows, using **Table across** for **Calculate as**, gives you the following results.

Customer Region	Service Line		
	Marketing	HR	Billing
APAC	3	2	1
EMEA	3	2	1
US	3	1	2

Percentile

The **Percentile** function calculates the percent of the values of the cells included in the calculation that are at or below the value for the given cell.

This percent is calculated as follows.

$$\text{percentile rank}(x) = 100 * B / N$$

Where:

B = number of scores below x

N = number of scores

Applying the **Percentile** function across the table rows, using **Table across** for **Calculate as**, gives you the following results.

Sum of Billed Amount by Customer Region, Date, Consumption Channel, and Service Line							
		Service Line > Consumption Channel					
Customer Region	Date	Billing			HR		
		Web	Mobile	API	Web	Mobile	API
APAC	2014	88.89	66.67	33.33	77.78	55.56	22.22
	2015	88.89	66.67	33.33	77.78	55.56	22.22
	2016	88.89	66.67	33.33	77.78	55.56	11.11
EMEA	2014	88.89	66.67	22.22	77.78	55.56	11.11
	2015	88.89	66.67	22.22	77.78	55.56	11.11
	2016	88.89	66.67	22.22	77.78	55.56	11.11
US	2014	77.78	55.56	11.11	88.89	66.67	22.22
	2015	77.78	55.56	22.22	88.89	66.67	11.11
	2016	77.78	55.56	11.11	88.89	66.67	22.22

Ways to apply pivot table calculations

You can apply table calculations in the ways described following. Table calculations are applied to only one field at a time. Thus, if you have a pivot table with multiple values, calculations are only applied to the cells representing the field that you applied the calculation to.

Topics

- [Table across](#)
- [Table down](#)
- [Table across down](#)
- [Table down across](#)
- [Group across](#)
- [Group down](#)
- [Group across down](#)
- [Group down across](#)

Table across

Using **Table across** applies the calculation across the rows of the pivot table, regardless of any grouping. This application is the default. For example, take the following pivot table.

		Consumption Channel		
Customer Region	Date	Web	Mobile	API
APAC	2014	\$599K	\$410K	\$186K
	2015	\$1,137K	\$795K	\$361K
	2016	\$1,997K	\$1,379K	\$638K
EMEA	2014	\$932K	\$646K	\$302K
	2015	\$1,511K	\$1,046K	\$531K
	2016	\$2,461K	\$1,702K	\$848K
US	2014	\$1,816K	\$1,257K	\$539K
	2015	\$3,230K	\$2,272K	\$963K
	2016	\$5,613K	\$3,971K	\$1,669K

Applying the **Running total** function using **Table across** gives you the following results, with row totals in the last column.

Sum of Billed Amount by Customer Region, Date, and Consumption Channel

		Consumption Channel		
Customer Region	Date	Web	Mobile	API
APAC	2014	\$599K	\$1,009K	\$1,195K
	2015	\$1,137K	\$1,932K	\$2,293K
	2016	\$1,997K	\$3,376K	\$4,014K
EMEA	2014	\$932K	\$1,578K	\$1,880K
	2015	\$1,511K	\$2,557K	\$3,088K
	2016	\$2,461K	\$4,163K	\$5,011K
US	2014	\$1,816K	\$3,074K	\$3,613K
	2015	\$3,230K	\$5,502K	\$6,465K
	2016	\$5,613K	\$9,584K	\$11,253K

Table down

Using **Table down** applies the calculation down the columns of the pivot table, regardless of any grouping.

Applying the **Running total** function using **Table down** gives you the following results, with column totals in the last row.

Sum of Billed Amount by Customer Region, Date, and Consumption Channel

Customer Region	Date	Consumption Channel		
		Web	Mobile	API
APAC	2014	\$599K	\$410K	\$186K
	2015	\$1,735K	\$1,205K	\$547K
	2016	\$3,733K	\$2,584K	\$1,185K
EMEA	2014	\$4,665K	\$3,230K	\$1,487K
	2015	\$6,176K	\$4,276K	\$2,017K
	2016	\$8,637K	\$5,978K	\$2,865K
US	2014	\$10,454K	\$7,235K	\$3,404K
	2015	\$13,684K	\$9,507K	\$4,367K
	2016	\$19,297K	\$13,478K	\$6,036K

Table across down

Using **Table across down** applies the calculation across the rows of the pivot table, and then takes the results and reapplies the calculation down the columns of the pivot table.

Applying the **Running total** function using **Table across down** gives you the following results. In this case, totals are summed both down and across, with the grand total in the lower-right cell.

Sum of Billed Amount by Customer Region, Date, and Consumption Channel

		Consumption Channel		
Customer Region	Date	Web	Mobile	API
APAC	2014	\$599K	\$1,009K	\$1,195K
	2015	\$2,332K	\$3,127K	\$3,488K
	2016	\$5,485K	\$6,864K	\$7,501K
EMEA	2014	\$8,433K	\$9,079K	\$9,381K
	2015	\$10,893K	\$11,939K	\$12,469K
	2016	\$14,931K	\$16,633K	\$17,480K
US	2014	\$19,297K	\$20,554K	\$21,093K
	2015	\$24,323K	\$26,595K	\$27,558K
	2016	\$33,171K	\$37,142K	\$38,811K

In this case, suppose that you apply the **Rank** function using **Table across down**. Doing so means that the initial ranks are determined across the table rows and then those ranks are in turn ranked down the columns. This approach gives you the following results.

Sum of Billed Amount by Customer Region, Date, and Consumption Channel				
Customer Region	Date	Consumption Channel		
		Web	Mobile	API
APAC	2014	21	24	27
	2015	13	18	25
	2016	6	11	20
EMEA	2014	16	19	26
	2015	10	14	23
	2016	4	8	17
US	2014	7	12	22
	2015	3	5	15
	2016	1	2	9

Table down across

Using **Table down across** applies the calculation down the columns of the pivot table. It then takes the results and reapplies the calculation across the rows of the pivot table.

You can apply the **Running total** function using **Table down across** to get the following results. In this case, totals are summed both down and across, with the grand total in the lower-right cell.

Sum of Billed Amount by Customer Region, Date, and Consumption Channel

		Consumption Channel		
Customer Region	Date	Web	Mobile	API
APAC	2014	\$599K	\$1,009K	\$1,195K
	2015	\$2,332K	\$3,127K	\$3,488K
	2016	\$5,485K	\$6,864K	\$7,501K
EMEA	2014	\$8,433K	\$9,079K	\$9,381K
	2015	\$10,893K	\$11,939K	\$12,469K
	2016	\$14,931K	\$16,633K	\$17,480K
US	2014	\$19,297K	\$20,554K	\$21,093K
	2015	\$24,323K	\$26,595K	\$27,558K
	2016	\$33,171K	\$37,142K	\$38,811K

You can apply the **Rank** function using **Table down across** to get the following results. In this case, the initial ranks are determined down the table columns. Then those ranks are in turn ranked across the rows.

Sum of Billed Amount by Customer Region, Date, and Consumption Channel				
		Consumption Channel		
Customer Region	Date	Web	Mobile	API
APAC	2014	21	24	27
	2015	13	18	25
	2016	6	11	20
EMEA	2014	16	19	26
	2015	10	14	23
	2016	4	8	17
US	2014	7	12	22
	2015	3	5	15
	2016	1	2	9

Group across

Using **Group across** applies the calculation across the rows of the pivot table within group boundaries, as determined by the second level of grouping applied to the columns. For example, if you group by field-2 and then by field-1, grouping is applied at the field-2 level. If you group by field-3, field-2, and field-1, grouping is again applied at the field-2 level. When there is no grouping, **Group across** returns the same results as **Table across**.

For example, take the following pivot table where columns are grouped by Service Line and then by Consumption Channel.

Sum of Billed Amount by Customer Region, Date, and Consumption Channel				
		Consumption Channel		
Customer Region	Date	Web	Mobile	API
APAC	2014	\$599K	\$410K	\$186K
	2015	\$1,137K	\$795K	\$361K
	2016	\$1,997K	\$1,379K	\$638K
EMEA	2014	\$932K	\$646K	\$302K
	2015	\$1,511K	\$1,046K	\$531K
	2016	\$2,461K	\$1,702K	\$848K
US	2014	\$1,816K	\$1,257K	\$539K
	2015	\$3,230K	\$2,272K	\$963K
	2016	\$5,613K	\$3,971K	\$1,669K

You can apply the **Running total** function using **Group across** to get the following results. In this case, the function is applied across the rows, bounded by the columns for each service category group. The Mobile columns display the total for both Consumption Channel values for the given Service Line, for the Customer Region and Date (year) represented by the given row. For example, the highlighted cell represents the total for the APAC region for 2012, for all Consumption Channel values in the Service Line named Billing.

Sum of Billed Amount by Customer Region, Date, Service Line, and Consumption Channel

Service Line > Consumption Channel

Customer Region	Date	Billing		Marketing	
		API	Mobile	API	Mobile
APAC	2012	\$23K	\$74K		
	2013	\$45K	\$153K	\$3K	\$11K
	2014	\$82K	\$260K	\$31K	\$100K
EMEA	2012	\$22K	\$78K		
	2013	\$50K	\$172K	\$4K	\$15K
	2014	\$122K	\$381K	\$66K	\$204K
US	2012	\$43K	\$140K		
	2013	\$114K	\$382K	\$14K	\$48K
	2014	\$213K	\$709K	\$110K	\$370K

Group down

Using **Group down** applies the calculation down the columns of the pivot table within group boundaries, as determined by the second level of grouping applied to the rows. For example, if you group by field-2 and then by field-1, grouping is applied at the field-2 level. If you group by field-3, field-2, and field-1, grouping is again applied at the field-2 level. When there is no grouping, **Group down** returns the same results as **Table down**.

For example, take the following pivot table where rows are grouped by Customer Region and then by Date (year).

Sum of Billed Amount by Customer Region, Date, and Consumption Channel







Customer Region	Date	Consumption Channel		
		Web	Mobile	API
APAC	2014	\$599K	\$410K	\$186K
	2015	\$1,137K	\$795K	\$361K
	2016	\$1,997K	\$1,379K	\$638K
EMEA	2014	\$932K	\$646K	\$302K
	2015	\$1,511K	\$1,046K	\$531K
	2016	\$2,461K	\$1,702K	\$848K
US	2014	\$1,816K	\$1,257K	\$539K
	2015	\$3,230K	\$2,272K	\$963K
	2016	\$5,613K	\$3,971K	\$1,669K

You can apply the **Running total** function using **Group down** to get the following results. In this case, the function is applied down the columns, bounded by the rows for each Customer Region group. The 2014 rows display the total for all years for the given Customer Region, for the Service Line and Consumption Channel represented by the given column. For example, the highlighted cell represents the total the APAC region, for the Billing service for the Mobile channel, for all the Date values (years) that display in the report.

Sum of Billed Amount by Customer Region, Date, Service Line, and Consumption Channel

Customer Region	Date	Service Line > Consumption Channel			
		Billing		Marketing	
		API	Mobile	API	Mobile
APAC	2012	\$23K	\$51K		
	2013	\$68K	\$159K	\$3K	\$8K
	2014	\$149K	\$337K	\$34K	\$77K
EMEA	2012	\$22K	\$56K		
	2013	\$72K	\$178K	\$4K	\$11K
	2014	\$194K	\$438K	\$70K	\$148K
US	2012	\$43K	\$97K		
	2013	\$157K	\$365K	\$14K	\$34K
	2014	\$370K	\$861K	\$124K	\$294K

Group across down

Using **Group across down** applies the calculation across the rows within group boundaries, as determined by the second level of grouping applied to the columns. Then the function takes the results and reapplies the calculation down the columns of the pivot table. It does so within group boundaries as determined by the second level of grouping applied to the rows.

For example, if you group a row or column by field-2 and then by field-1, grouping is applied at the field-2 level. If you group by field-3, field-2, and field-1, grouping is again applied at the field-2 level. When there is no grouping, **Group across down** returns the same results as **Table across down**.

For example, take the following pivot table where columns are grouped by Service Line and then by Consumption Channel. Rows are grouped by Customer Region and then by Date (year).

Sum of Billed Amount by Customer Region, Date, and Consumption Channel				
		Consumption Channel		
Customer Region	Date	Web	Mobile	API
APAC	2014	\$599K	\$410K	\$186K
	2015	\$1,137K	\$795K	\$361K
	2016	\$1,997K	\$1,379K	\$638K
EMEA	2014	\$932K	\$646K	\$302K
	2015	\$1,511K	\$1,046K	\$531K
	2016	\$2,461K	\$1,702K	\$848K
US	2014	\$1,816K	\$1,257K	\$539K
	2015	\$3,230K	\$2,272K	\$963K
	2016	\$5,613K	\$3,971K	\$1,669K

You can apply the **Running total** function using **Group across down** to get the following results. In this case, totals are summed both down and across within the group boundaries. Here, these boundaries are Service Line for the columns and Customer Region for the rows. The grand total appears in the lower-right cell for the group.

Sum of Billed Amount by Customer Region, Date, Service Line, and Consumption Channel

		Service Line > Consumption Channel			
Customer Region	Date	Billing		Marketing	
		API	Mobile	API	Mobile
APAC	2012	\$23K	\$201K		
	2013	\$68K	\$309K	\$3K	\$42K
	2014	\$149K	\$487K	\$34K	\$111K
EMEA	2012	\$22K	\$250K		
	2013	\$72K	\$372K	\$4K	\$81K
	2014	\$194K	\$632K	\$70K	\$219K
US	2012	\$43K	\$466K		
	2013	\$157K	\$734K	\$14K	\$157K
	2014	\$370K	\$1,231K	\$124K	\$418K

You can apply the **Rank** function using **Group across down** to get the following results. In this case, the function is first applied across the rows bounded by each Service Line group. The function is then applied again to the results of that first calculation, this time applied down the columns bounded by each Customer Region group.

Sum of Billed Amount by Customer Region, Date, Service Line, and Consumption Channel

		Service Line > Consumption Channel			
Customer Region	Date	Billing		Marketing	
		API	Mobile	API	Mobile
APAC	2012	6	4		
	2013	5	2	4	3
	2014	3	1	2	1
EMEA	2012	6	4		
	2013	5	2	4	3
	2014	3	1	2	1
US	2012	6	5		
	2013	4	2	4	3
	2014	3	1	2	1

Group down across

Using **Group down across** applies a calculation down the columns within group boundaries, as determined by the second level of grouping applied to the rows. Then Amazon Quick takes the results and reapplies the calculation across the rows of the pivot table. Again, it reapplies the calculation within group boundaries as determined by the second level of grouping applied to the columns.

For example, if you group a row or column by field-2 and then by field-1, grouping is applied at the field-2 level. If you group by field-3, field-2, and field-1, grouping is again applied at the field-2 level. When there is no grouping, **Group down across** returns the same results as **Table down across**.

For example, take the following pivot table. Columns are grouped by Service Line and then by Consumption Channel. Rows are grouped by Customer Region and then by Date (year).

		Consumption Channel		
Customer Region	Date	Web	Mobile	API
APAC	2014	\$599K	\$410K	\$186K
	2015	\$1,137K	\$795K	\$361K
	2016	\$1,997K	\$1,379K	\$638K
EMEA	2014	\$932K	\$646K	\$302K
	2015	\$1,511K	\$1,046K	\$531K
	2016	\$2,461K	\$1,702K	\$848K
US	2014	\$1,816K	\$1,257K	\$539K
	2015	\$3,230K	\$2,272K	\$963K
	2016	\$5,613K	\$3,971K	\$1,669K

You can apply the **Running total** function using **Group down across** to get the following results. In this case, totals are summed both down and across within the group boundaries. In this case, these are Service Category for the columns and Customer Region for the rows. The grand total is in the lower-right cell for the group.

Sum of Billed Amount by Customer Region, Date, Service Line, and Consumption Channel

		Service Line > Consumption Channel			
Customer Region	Date	Billing		Marketing	
		API	Mobile	API	Mobile
APAC	2012	\$23K	\$201K		
	2013	\$68K	\$309K	\$3K	\$42K
	2014	\$149K	\$487K	\$34K	\$111K
EMEA	2012	\$22K	\$250K		
	2013	\$72K	\$372K	\$4K	\$81K
	2014	\$194K	\$632K	\$70K	\$219K
US	2012	\$43K	\$466K		
	2013	\$157K	\$734K	\$14K	\$157K
	2014	\$370K	\$1,231K	\$124K	\$418K

You can apply the **Rank** function using **Group down across** to get the following results. In this case, the function is first applied down the columns bounded by each Customer Region group. The function is then applied again to the results of that first calculation, this time applied across the rows bounded by each Service Line group.

Sum of Billed Amount by Customer Region, Date, Service Line, and Consumption Channel

		Service Line > Consumption Channel			
Customer Region	Date	Billing		Marketing	
		API	Mobile	API	Mobile
APAC	2012	6	4		
	2013	5	2	4	3
	2014	3	1	2	1
EMEA	2012	6	4		
	2013	5	2	4	3
	2014	3	1	2	1
US	2012	6	5		
	2013	4	2	4	3
	2014	3	1	2	1

Pivot table limitations

The following limitations apply to pivot tables:

- You can create pivot tables with up to 500,000 records.
- You can add any combination of row and column field values that add up to 40. For example, if you have 10 row field values, then you can add up to 30 column field values.
- You can create pivot table calculations only on nonaggregated values. For example, if you create a calculated field that is a sum of a measure, you can't also add a pivot table calculation to it.
- If you are sorting by a custom metric, you can't add a table calculation until you remove the custom metric sort.
- If you are using a table calculation and then add a custom metric, you can't sort by the custom metric.
- Totals and subtotals are blank for table calculations on metrics aggregated by distinct count.

Pivot table best practices

It's best to deploy a minimal set of rows, columns, metrics, and table calculations, rather than offering all possible combinations in one pivot table. If you include too many, you risk overwhelming the viewer and you can also run into the computational limitations of the underlying database.

To reduce the level of complexity and reduce the potential for errors, you can take the following actions:

- Apply filters to reduce the data included in for the visual.
- Use fewer fields in the **Row** and **Column** field wells.
- Use as few fields as possible in the **Values** field well.
- Create additional pivot tables so that each displays fewer metrics.

In some cases, there's a business need to examine many metrics in relation to each other. In these cases, it can be better to use multiple visuals on the same dashboard, each showing a single metric. You can reduce the size of the visuals on the dashboard, and colocate them to form a grouping. If a decision the viewer makes based on one visual creates the need for a different view, you can deploy custom URL actions to launch another dashboard according to the choices made by the user.


It's best to think of visuals as building blocks. Rather than using one visual for multiple purposes, use each visual to facilitate one aspect of a larger business decision. The viewer should have enough data to make a well-informed decision, without being overwhelmed by the inclusion of all possibilities.

Using radar charts

You can use radar charts, which are also known as spider charts, to visualize multivariate data in Amazon Quick. In a radar chart, one or more groups of values are plotted over multiple common variables. Each variable has its own axis, and each axis is arranged radially around a central point. The data points from a single observation are plotted on each axis and connected to each other to form a polygon. Multiple observations can be plotted in a single radar chart to display multiple polygons, which makes it easier to spot outlying values for multiple observations quickly.

In Quick, you can organize a radar chart along its category, value, or color axes by dragging and dropping fields to the **Category**, **Value**, and **Color** field wells. How you choose to distribute fields among the field wells determines the axis that the data is plotted on.

The following image shows an example of a radar chart.



Radar chart features

To view the features supported by radar charts, use the following table.

Feature	Supported?	Comments	For more information
Changing the legend display	Yes		Legends on visual types in Quick

Feature	Supported?	Comments	For more information
Changing the title display	Yes		Titles and subtitles on visual types in Quick
Changing the axis range	Yes		Range and scale on visual types in Quick
Changing the visual colors	Yes		Colors in visual types in Quick
Focusing on or excluding elements	Yes		Focusing on visual elements Excluding visual elements
Sorting	Limited	You can only sort data fields that are in the Category and Color field wells.	Sorting visual data in Amazon Quick
Performing field aggregation	Yes		Changing field aggregation
Adding drill-downs	Not supported		Adding drill-downs to visual data in Quick Sight
Choosing size	Yes		Formatting in Amazon Quick
Showing totals	Not supported		Formatting in Amazon Quick

Creating a radar chart

Use the following procedure to create a radar chart.

To create a radar chart

1. On the analysis page, choose **Visualize** on the tool bar.
2. Choose **Add** on the application bar, and then choose **Add visual**.

3. On the **Visual types** pane, choose the radar chart icon.
4. From the **Fields list** pane, drag the fields that you want to use to the appropriate field wells. In most cases, you want to use dimension or measure fields as indicated by the target field well.

To create a radar chart, drag fields to the **Category**, **Value**, and **Group/Color** field wells. The axis that a radar chart is organized around is determined by the way that you organize fields into their respective field wells:

- In a radar chart that uses a **value axis**, dimension values are shown as lines and axes represent value fields. To create a radar chart that uses a value axis, add one category field to the **Color** field well and one or more values to the **Value** field well.
- In a radar chart that uses a **dimension axis**, group dimension values are shown as axes and value fields are shown as lines. All axes share a range and scale. To create a radar chart that uses a dimension axis, add one dimension to the **Group** field well and one or more values to the **Value** field well.
- In a radar chart that uses a **dimension-color axis**, group dimension values are shown as axes and color dimension values are shown as lines. All axes share a range and scale. To create a radar chart that uses a dimension-color axis, add one dimension to the **Category** field well, one value to the **Value** field well, and one dimension to the **Color** field well.

Using Sankey diagrams

Use Sankey diagrams to show flows from one category to another, or paths from one stage to the next.

For example, a Sankey diagram can show the number of people migrating from one country to another. A Sankey diagram can also show the path a web visitor takes from one page to the next on a company website, with possible stops along the way.

Data for Sankey diagrams

To create Sankey diagrams in Quick, your dataset should contain a measure and two dimensions (one dimension containing source categories and another containing destination categories).

The following table is a simple example of data for a Sankey diagram.

Dimension (Source)	Dimension (Destination)	Measure (Weight)
--------------------	-------------------------	------------------

Dimension (Source)	Dimension (Destination)	Measure (Weight)
A	W	500
A	X	23
A	Y	147

The following Sankey diagram is created when the dimensions and measure are added to the field well, with the A node on the left linking to the W, Y, and X nodes on the right. The width of each link between nodes is determined by the value in the Measure (Weight) column. The nodes are automatically ordered.

To create multilevel Sankey diagrams in Amazon Quick, your dataset should still contain a measure and two dimensions (one for source and one for destination), but in this case your data values differ.

The following table is a simple example of data for a multilevel Sankey diagram with two stages.

Dimension (Source)	Dimension (Destination)	Measure (Weight)
A	W	500
A	X	23
A	Y	147
W	Z	300
X	Z	5
Y	Z	50

The following Sankey diagram is created when the dimensions and measure are added to the field well. Here, the A node on the left links to the W, Y, and X nodes in the middle, and the W, Y, and X nodes then link to the Z node on the right. The width of each link between nodes is determined by the value in the Measure (Weight) column.

Working with cyclical data

Sometimes, the data that you use for a Sankey diagram contains cycles. For example, suppose that you're visualizing user traffic flows between pages on a website. You might discover that users who come to page A move to page E, and then come back to page A. An entire flow might look something like A-E-A-B-A-E-A.

When your data contains cycles, the nodes in each cycle are repeated in Quick. For example, if your data contains the flow A-E-A-B-A-E-A, the following Sankey diagram is created.

Preparing data for Sankey diagrams

If your dataset doesn't contain Source or Destination columns, prepare your data to include them. You can prepare data when creating a new dataset, or when editing an existing dataset. For more information about creating a new dataset and preparing it, see [Creating datasets](#). For more information about opening an existing dataset for data preparation, see [Editing datasets](#).

The following procedure uses an example table (illustrated in following) to demonstrate how to prepare your data for Sankey diagrams in Quick. The table includes three columns: Customer ID, Time, and Action.

Customer ID	Time	Action
1	9:05 am	Step 1
1	9:06 am	Step 2
1	9:08 am	Step 3
2	11:44 am	Step 1
2	11:47 am	Step 2

Customer ID	Time	Action
2	11:48 am	Step 3

To create a Sankey diagram in Quick using this data, first add Source and Destination columns to the table. Use the following procedure to learn how.

To add Source and Destination columns to your table

1. Add a Step Number column to the table to number or rank each row.

There are multiple ways to compute the Step Number column. If your data source is compatible with SQL and your database supports ROW_NUMBER or RANK functions, you can use custom SQL in Quick to order the rows in the Step Number column. For more information about using custom SQL in Quick, see [Using SQL to customize data](#).

Customer ID	Time	Action	Step Number
1	9:05 am	Step 1	1
1	9:06 am	Step 2	2
1	9:08 am	Step 3	3
2	11:44 am	Step 1	1
2	11:47 am	Step 2	2
2	11:48 am	Step 3	3

2. Add a Next Row Number column to the table with values equal to Step Number plus one.

For example, in the first data row of the table, the value for Step Number is 1. To compute the value for Next Step Number for that row, add 1 to that value.

$$1 + 1 = 2$$

The value for Step Number in the second data row of the table is 2; therefore, the value for Next Step Number is 3.

$$2 + 1 = 3$$

Customer ID	Time	Action	Step Number	Next Step Number
1	9:05 am	Step 1	1	2
1	9:06 am	Step 2	2	3
1	9:08 am	Step 3	3	4
2	11:44 am	Step 1	1	2
2	11:47 am	Step 2	2	3
2	11:48 am	Step 3	3	4

3. Join the table with itself:
 - a. For **Join type**, choose **Inner**.
 - b. For **Join clauses**, do the following:
 - i. Choose **Customer ID = Customer ID**
 - ii. Choose **Next Step Number = Step Number**

Joining the two tables creates two columns for Customer ID, Time, Action, Step Number and Next Step Number. The columns from the table at the left of the join are Source columns. The columns from the table at the right of the join are Destination columns.

For more information about joining data in Quick, see [Joining data](#).

4. (Optional) Rename columns to indicate sources and destinations.

The following is an example:

1. Rename the **Action** column on the left to **Source**.
2. Rename the **Action [copy]** column on the right to **Destination**.
3. Rename the **Time** column on the left to **Start Time**.
4. Rename the **Time [copy]** column on the right to **End Time**.

Your data is now ready to visualize.

Creating Sankey diagrams

Use the following procedure to create a Sankey diagram.

To create a Sankey diagram

1. On the analysis screen, choose **Visualize** on the left toolbar.
2. On the application bar, choose **Add**, and then choose **Add visual**.
3. On the **Visual types** pane, choose the Sankey diagram icon.
4. On the menu in the upper-right corner of the visual, choose the **Properties** icon.
5. In the **Properties** pane, choose either the **Source** or **Destination** section.

Customizing the number of nodes

Use the following procedure to customize the number of nodes that appear in a Sankey diagram. Quick supports up to 100 Source/Destination nodes.

To customize the number of nodes that appear in a Sankey diagram

1. On the analysis page, choose the Sankey diagram visual that you want to format.
2. On the menu in the upper-right corner of the visual, select the **Format Visual** icon.
3. In the **Properties** pane that opens, choose either the **Source** or **Destination** tab.
4. For **Number of nodes displayed**, enter a number.

The nodes in the diagram update to the number that you specified. The top nodes are automatically shown. All other nodes are placed in an **Other** category.



Note

Specifying the number of Source nodes controls how many Source nodes can appear overall in the diagram. Specifying the number of Destination nodes controls how many Destination nodes can appear per Source node. This means that if there is more than one Source node in your diagram, the overall number of Destination nodes will be higher than the number specified.

Quick supports up to 100 Source/Destination nodes.

For example, the following Sankey diagram has a limit of three source nodes (out of five), so the top three are shown in the diagram. The other two source nodes are placed in the Other category.

To remove the **Other** category from the diagram, select it in the view and choose **Hide “other” categories**.

Sankey diagram features

To understand the features supported by Sankey diagrams, use the following table.

Feature	Supported?	For more information	
Changing the legend display	No		
Changing the title display	Yes	Titles and subtitles on visual types in Quick	
Changing the axis range	No		
Changing the visual colors	No		
Focusing on or excluding elements	Yes	Focusing on visual elements	

Feature	Supported?	For more information	
		Excluding visual elements	
Sorting	No		
Performing field aggregation	Yes	Changing field aggregation	
Adding drill-downs	No		
Conditional formatting	No		

Using scatter plots

Use scatter plots to visualize two or three measures across two dimensions.

Each bubble on the scatter plot represents one or two dimension values. The X and Y axes represent two different measures that apply to the dimension. A bubble appears on the chart at the point where the values for the two measures for an item in the dimension intersect. Optionally, you can also use bubble size to represent an additional measure.

Scatter plots show up to 2500 datapoints in aggregated and unaggregated scenarios regardless of whether a color or label dimension is used in the visual. Due to the order of limit operations, there may be cases where fewer datapoints for a dataset are shown. For more information about how Amazon Quick handles data that falls outside display limits, see [Display limits](#).

Scatter plot features

To understand the features supported by scatter plots, use the following table.

Feature	Supported?	Comments	For more information
Changing the legend display	Yes, with exceptions	Scatter plots display a legend if you have the	Legends on visual types in Quick

Feature	Supported?	Comments	For more information
		Group/Color field well populated.	
Changing the title display	Yes		Titles and subtitles on visual types in Quick
Changing the axis range	Yes	You can set the range for both the X and Y axes.	Range and scale on visual types in Quick
Showing or hiding axis lines, grid lines, axis labels, and axis sort icons	Yes		Axes and grid lines on visual types in Quick
Changing the visual colors	Yes		Colors in visual types in Quick
Focusing on or excluding elements	Yes, with exceptions	You can focus on or exclude a bubble in a scatter plot, except when you are using a date field as a dimension . In that case, you can only focus on a bubble, not exclude it.	Focusing on visual elements Excluding visual elements
Sorting	No		Sorting visual data in Amazon Quick
Performing field aggregation	Yes	You must apply aggregation to the fields you choose for the X axis, Y axis, and size, and can't apply aggregation to the field that you choose for the group or color.	Changing field aggregation

Feature	Supported?	Comments	For more information
Displaying unaggregated fields	Yes	On the field context menu, choose None to display unaggregated X and Y axis values. If your scatter plot shows unaggregated fields, you can't apply aggregations to the field that is in the color or label field well. Mixed aggregation is not supported for scatter plots.	
Adding drill-downs	Yes	You can add drill-down levels to the Group/Color field well.	Adding drill-downs to visual data in Quick Sight

Creating a scatter plot

Use the following procedure to create a scatter plot.

To create a scatter plot

1. On the analysis page, choose **Visualize** on the tool bar.
2. Choose **Add** on the application bar, and then choose **Add visual**.
3. On the **Visual types** pane, choose the scatter plot icon.
4. From the **Fields list** pane, drag the fields that you want to use to the appropriate field wells. Typically, you want to use dimension or measure fields as indicated by the target field well. If you choose to use a dimension field as a measure, the **Count** aggregate function is automatically applied to it to create a numeric value.

To create a scatter plot, drag a measure to the **X axis** field well, a measure to the **Y axis** field well, and a dimension to the **Color** or **Label** field well. To represent another measure with bubble size, drag that measure to the **Size** field well.

5. (Optional) Add drill-down layers by dragging one or more additional fields to the **Color** field well. For more information about adding drill-downs, see [Adding drill-downs to visual data in Quick Sight](#).


Scatter plot use cases

You can choose to plot unaggregated values even if you are using a field on Color by using the aggregate option **none** on the field menu, which also contains aggregation options like **sum**, **min**, and **max**. If one value is set to be aggregated, the other value will be automatically set as aggregated. The same applies to unaggregated scenarios. Mixed aggregation scenarios are not supported, meaning that one value cannot be set as aggregated while the other is unaggregated. Note that the unaggregated scenario, which is the **none** option, is supported only for numerical values, while categorical values, such as dates or dimensions, will display only aggregate values, such as **count** and **count distinct**.


Using the **none** option, you can choose to set both X and Y values to either aggregated or unaggregated from the **X axis** and **Y axis** field menus. This will define whether or not values will be aggregated by dimensions in the **Color** and **Label** field wells. To get started, add the required fields and choose the appropriate aggregation based on your use case, as shown in the following sections.

Unaggregated use cases

- Unaggregated X and Y values with Color



- Unaggregated X and Y values with Label




• Unaggregated X and Y values with Color and Label



Aggregated use cases


- Aggregated X and Y values with Color



- Aggregated X and Y values with Label



• Aggregated X and Y values with Color and Label



Using tables as visuals

Use a table visual to see a customized table view of your data. To create a table visual, choose at least one field of any data type. You can add as many columns as you need, up to 200. You can also add calculated columns.

Table visuals don't display a legend. You can hide or display the title on a table. You can also hide or display totals, and choose to show totals at the top or the bottom of the table. For more information, see [Analytics formatting per type in Quick](#).

To create a table visual

1. Open Amazon Quick and choose **Analyses** on the navigation pane at left.
2. Choose one of the following:
 - To create a new analysis, choose **New analysis** at upper right. For more information, see [Starting an analysis in Quick Sight](#).
 - To use an existing analysis, choose the analysis that you want to edit.
3. Choose **Insert** from the file menu and then **Add Visual**.
4. At lower left, choose the table icon from **Visual types**.
5. On the **Fields** list pane, choose the fields that you want to use. If you want to add a calculated field, choose **Insert** on the file menu and then **Add Calculated Field**.

To create a nonaggregated view of the data, add fields only to the **Value** field well. Doing this shows data without any aggregations.

To create an aggregated view of the data, choose the fields that you want to aggregate by, and then add them to the **Group by** field well.

To show or hide columns on a table

1. On your visual, choose the field that you want to hide, then choose **Hide column**.
2. To display hidden columns, choose any column, then choose **Show all hidden columns**.

To transpose columns to rows and rows to columns

- Choose the transpose icon (




)

near the top right of the visual. It has two arrows at a 90 degree angle.

To vertically align columns

- On your visual, choose the **Format visual** icon (




)

near the top right of the visual.

- In the **Properties** pane, choose **Table options**, and choose your table's vertical alignment.

To wrap the text for headers

- On your visual, choose the **Format visual** icon (



)

near the top right of the visual.

- In the **Properties** pane, choose **Table options**, and select **Wrap header text**.

To rearrange columns in a table chart

- Open the analysis with the visual that you want to sort. Visuals pane will be open by default.
- Do one of the following:
 - Drag and drop one or more fields in **Field wells** to rearrange their order.
 - Select a field directly in the table and choose the left or right arrow on **Move column**.

Using field styling

You can render URLs in a table as links by using the **Field styling** pane of the format visual menu. You can add up to 500 rows of links for each page in a table. Only https and mailto hyperlinks are supported.

To add links to your tables

1. From the Quick homepage, choose **Analyses**, and then choose the analysis that you want to customize.
2. Choose the table that you want to change.
3. On the menu at the upper right of the table, choose **Format visual**.
4. For **Format visual**, choose **Field styling**.
5. On the **Field styling** pane, choose the field that you want to style from the menu.
6. In the **Url options** section of the **Field styling** menu, choose **Make URLs hyperlinks**.

After you add links to your table, you can choose where you want the links to open when they're selected in the **Open in** section of the **Field style** pane. You can choose to have links open in a new tab, a new window, or in the same tab.

You can also choose how you want to style the link in the **Style as** section of **Field style** pane. Your links can appear as hyperlinks, icons, or plain text, or you can set a custom link.

To adjust the font size of a link icon or URL, change the **Font size** in the **Cells** section of the **Table options** pane of the **Format visual** menu.

You can set any URLs in your table that point to images to render in the table as images. Doing this can be useful when you want to include an image of a product as a part of a table.

To show URLs as images

1. From the Quick home page, choose **Analyses**, and then choose the analysis that you want to customize.
2. Choose the table that you want to change.
3. On the menu at the upper-right of the table, choose **Format visual**.
4. In the **Format visual** menu, choose **Field styling**.
5. In the **Field styling** pane, choose the field that you want to style from the menu.
6. In the **Url options** section of the **Field styling** menu, choose **Show URLs as images**.

After rendering images in a table, you can choose how to size the images in the **Image sizing** section of the **Field style** pane. You can fit images to their cell's height or width, or you can choose not to scale the image. Images fit to a cell's height by default.

Freeze columns to table visuals

You can freeze columns to your table visuals to lock specific columns in place on screen. This allows essential information to remain visible while readers scroll across the table. You can freeze columns one at a time, or you can freeze groups of columns in one action. All pinned columns are fixed to the far left side of the table and stay visible on screen at all times. This allows Quick readers to have a constant reference point for key data or information as they interact with other parts of the table.

To freeze columns to a table

1. On the table that you want to freeze a column to, choose the column that you want to pin.
2. Choose one of the following options.
 - To freeze a single column, choose **Freeze column**.
 - To freeze all columns up to the column that you choose, choose **Freeze up to this column**.

If your table has multiple pinned columns, you can reorder the columns in the order that you want. To adjust the order of the pinned columns on a table, choose the header of the column that you want to move, and then choose **Move** in the direction that you want.

To unfreeze columns from a table

1. On the table that you want to change, choose the pinned column that you want to unpin.
2. Choose one of the following options.
 - To unfreeze a single column, choose **Unfreeze column**.
 - To unfreeze all frozen columns, choose **Unfreeze all columns**.

Custom total values

Quick authors can define the total and subtotal aggregations for their table or pivot table visuals from the field wells. For tables, the custom total menu is only available if totals are toggled on for the visual.

To change the aggregation of a total or subtotal

1. Navigate to the analysis that you want to change, and choose the table or pivot table visual whose total you want to define.


2. Choose the field that you want to change from the field wells.
3. Choose **Total**, and then choose the aggregation that you want. The following options are available.
 - **Default** – The total calculation uses the same aggregation as the metric field.
 - **Sum** – Calculates the sum of the data in the visual.
 - **Average** – Calculates the average of the data in the visual.
 - **Min** – Calculates the minimum value of the data in the visual.
 - **Max** – Calculates the maximum value of the data in the visual.
 - **None (HIDE)** – Totals are not calculated. When you choose this option, the total and subtotal cells in the visual are left blank. If the outer dimension is sorted with the metric field that calculates the total or subtotal, the dimension is sorted alphabetically. When you change the value from **None (HIDE)** to another value, the outer dimension is sorted by the subtotals that are calculated with the specified aggregation type.

The following limitations apply to custom totals.

- Conditional formatting is not supported for custom totals.
- Total aggregations aren't supported for string columns. Total aggregations include **Min**, **Max**, **Sum**, and **Average**.
- Date columns are incompatible with **Average** and **Sum** total aggregation functions.

Sorting tables

In Amazon Quick, you can sort values in a table by fields in the columns headers of the table or with the **Sort visual** tool. You can sort up to 10 columns in a single table. Quick can also use an off-visual sort. You can sort columns in an **Ascending** or a **Descending** order. The following image shows the **Sort visual** icon and pop over.



Single column sort options

Quick Authors can access single column sort options from the field wells, the column headers, or from the **Sort visual** menu. Use the procedure below to use set up a single column sort on a table in Quick.

1. Open the [Quick console](#).
2. Open the analysis that you want to work in and navigate to the table that you want to sort.
3. Choose the header of the column that you want to sort.
4. For **Sort by**, choose the arrow icon, and then choose the field that you want to sort by.

You can also set up a single column sort in the **Sort visual** menu. To access the sort visual menu, choose the **Sort visual** icon in the on-visual menu. In the **Sort visual** menu, choose the field that you want to sort by, and then choose if you want the sort in an ascending or descending order. By default, new sorts are sorted in an ascending order. When you are finished, choose **APPLY**.

Tables that use single column sorting are sorted one column at a time. When a user chooses a new column to sort by, the previous sort order is overridden.

To make changes to a single column sort, open the **Sort visual** menu annd use the dropdown menus to choose a new field or sort order. When you are finished with your changes, choose **APPLY**.

To reset a table to its original state, open the **Sort visual** menu and choose **RESET**.

Multi column sort options

Quick authors can access multi column sort options from the **Sort visual** menu. Use the procedure below to set up a multi column sort for a table.

1. Open the [Quick console](#).

2. Open the analysis that you want to work in and navigate to the table that you want to sort.
3. Choose the **Sort visual** icon to open the **Sort visual** menu.
 - a. Alternatively, choose a header that you want to sort.
 - b. For **Sort by**, choose the arrow icon, and then choose **Multiple fields**.
4. In the **Sort visual** menu that opens, choose a field from the **Sort by** dropdown, and then choose whether you want the field sorted in an ascending or descending order.
5. To add another sort, choose **ADD SORT**, and repeat the workflow from Step 4. You can add up to 10 sorts to each table.
6. When you are finished, choose **APPLY**.

Columns are sorted in the order that they are added to the **Sort visual** menu. To change the order that columns are sorted by, open the **Sort visual** menu and use the **Sort by** dropdowns to reorder the sorts. When you are finished, choose **APPLY** to apply the new sort order to the table.

To reset a table to its original state, open the **Sort visual** menu and choose **RESET**.

Off visual sort options

Quick authors can configure an off-visual sort to sort the values in a table by a field and aggregation that is a part of the dataset that the table uses but not in one of the table's field wells. One off-field sort can be configured to a single table at a time.

Use the procedure below to configure an off-visual sort.

To add an off-visual sort to a table

1. Open the [Quick console](#).
2. Open the analysis that you want to work in and navigate to the table that you want to sort.
3. Choose the header of any column in the table.
4. For **Sort by**, choose the arrow icon, and then choose **Off-visual field**.
5. In the **Off-visual field** pane that appears, open the **Sort by** dropdown menu and choose the field that you want to sort.
6. For **Aggregation** open the dropdown menu and choose the aggregation that you want to use.
7. For **Sort order**, choose if you want the sort to be in an ascending or descending order.
8. When you are finished, choose **Apply**.

After a off-visual sort is applied to a table, the sort is shown in the **Sort visual** menu. The sort order of a table that contains an off-visual sort depends on the sort configuration of the table when the off-visual sort is added. If an off-visual sort is added to a table that already has a single or multi column sort configured, the off-visual sort overrides all other sorts. If the off-visual sort is applied before single or multi column sorts, you can add and reorder more sorts to the table.

Using text boxes

Add text to add context to sheets in an analysis by using a text box. Text can hold directions, descriptions, or even hyperlinks to external websites. The toolbar on the text box offers font settings so you can customize the font type, style, color, size, spacing, size in pixels, text highlights, and alignment. The text box itself has no format settings.

To add text to a new text box, simply select it and begin typing.

Using tree maps

To visualize one or two measures for a dimension, use tree maps.

Each rectangle on the tree map represents one item in the dimension. Rectangle size represents the proportion of the value for the selected measure that the item represents compared to the whole for the dimension. You can optionally use rectangle color to represent another measure for the item. Rectangle color represents where the value for the item falls in the range for the measure, with darker colors indicating higher values and lighter colors indicating lower ones.

Tree maps show up to 100 data points for the **Group by** field. For more information about how Amazon Quick handles data that falls outside display limits, see [Display limits](#).

Tree map features

To understand the features supported by tree maps, use the following table.

Feature	Supported?	Comments	For more information
Changing the legend display	Yes		Legends on visual types in Quick
Changing the title display	Yes		Titles and subtitles on visual types in Quick

Feature	Supported?	Comments	For more information
Changing the axis range	Not applicable		Range and scale on visual types in Quick
Changing the visual colors	No		Colors in visual types in Quick
Focusing on or excluding elements	Yes, with exceptions	You can focus on or exclude a rectangle from a tree map, except when you are using a date field as the dimension. In that case, you can only focus on a rectangle, not exclude it.	Focusing on visual elements Excluding visual elements
Sorting	No	Default sorting is in descending order by the measure in the Size column.	Sorting visual data in Amazon Quick
Performing field aggregation	Yes	You must apply aggregation to the fields you choose for size and color, and can't apply aggregation to the field that you choose to group by.	Changing field aggregation
Adding drill-downs	Yes	You can add drill-down levels to the Group by field well.	Adding drill-downs to visual data in Quick Sight

Creating a tree map

Use the following procedure to create a tree map.

To create a tree map

1. On the analysis page, choose **Visualize** on the tool bar.

2. Choose **Add** on the application bar, and then choose **Add visual**.
3. On the **Visual types** pane, choose the tree map icon.
4. From the **Fields list** pane, drag the fields that you want to use to the appropriate field wells. Typically, you want to use dimension or measure fields as indicated by the target field well. If you choose to use a dimension field as a measure, the **Count** aggregate function is automatically applied to it to create a numeric value.

To create a tree map, drag a measure to the **Size** field well and a dimension to the **Group by** field well. Optionally, drag another measure to the **Color** field well.

5. (Optional) Add drill-down layers by dragging one or more additional fields to the **Group by** field well. For more information about adding drill-downs, see [Adding drill-downs to visual data in Quick Sight](#).


Using waterfall charts

Use a waterfall chart to visualize a sequential summation as values are added or subtracted. In a waterfall chart, the initial value goes through a (positive or negative) change, with each change represented as a bar. The final total is represented by the last bar. Waterfall charts are also known as *bridges* because the connectors between the bars bridge the bars together, showing that they visually belong to the same story.

Waterfall charts are most commonly used to present financial data, because you can show change within one time period or from one time period to another. This way, you can visualize the different factors that have an impact your project cost. For example, you can use a waterfall chart to show gross sales to net income within the same month, or the difference in net income from last year to this year, and the factors that were responsible for this change.

You can also use waterfall charts to present statistical data, for example how many new employees you hired and how many employees left your company within a year.

The following screenshot shows a waterfall chart.



To create a basic waterfall chart visual

1. Open Amazon Quick and choose **Analyses** on the navigation pane at left.
2. Choose one of the following:
 - To create a new analysis, choose **New analysis** at upper right. For more information, see [Starting an analysis in Quick Sight](#).
 - To use an existing analysis, choose the analysis that you want to edit.
3. Choose **Add (+), Add Visual**.
4. At lower left, choose the waterfall chart icon from **Visual types**.
5. On the **Fields list** pane, choose the fields that you want to use for the appropriate field wells. Waterfall charts require one category or measure in **Value**.
6. (Optional) Add drill-down layers by dragging one or more additional fields to the **Group/Color** field well. For more information about adding drill-downs, see [Adding drill-downs to visual data in Quick Sight](#).

To understand the features supported by waterfall charts, see [Analytics formatting per type in Quick](#). For customization options, see [Formatting in Amazon Quick](#).

Using word clouds

As an engaging way to display how often a word is used in relation to other words in a dataset, use word clouds. The best use for this type of visual is to show word or phrase frequency. It can also make a fun addition to show trending items or actions. You can use a fixed dataset for

Each word in a word cloud represents one or more values in a dimension. The size of the word represents the frequency of a value's occurrence in a selected dimension, in proportion to the occurrences of other values in the same dimension. Word clouds are best when precision isn't important and there aren't a large number of distinct values.

Sum of Weighted Revenue by Lead Name

SHOWING TOP 100 IN LEAD NAME

US_Enterprise_977

US_Enterprise_1402

US_Enterprise_167

US_Startup_1182

US_Startup_35

US_Enterprise_167

US_Startup_451

US_Startup_212

US_Enterprise_167

US_Startup_1108

US_Startup_475

US_Enterprise_167

US_Startup_966

APAC_Enterprise_148

US_Enterprise_821

EMEA_Enterprise_674

US_Startup_696

US_Startup_457

APAC_Enterprise_428

US_Startup_438

US_Enterprise_137

US_Startup_1164

APAC_Enterprise_702

US_Startup_267

US_Enterprise_62

APAC_Startup_128

US_Startup_762

US_Enterprise_1168

US_Enterprise_1074

US_Startup_963

APAC_Enterprise_401

APAC_Enterprise_55

APAC_Startup_80

US_Startup_1082

APAC_Enterprise_266

APAC_Enterprise_124

APAC_Startup_104

EMEA_Startup_110

US_Startup_197

APAC_Startup_105

APAC_Startup_69

EMEA_Startup_225

US_Enterprise_198

US_Startup_407

US_Startup_92

EMEA_Startup_214

US_Startup_721

APAC_Startup_106

APAC_Startup_21

US_Enterprise_259

US_Startup_28

APAC_Enterprise_390

APAC_Enterprise_172

APAC_Startup_32

US_Enterprise_1323

APAC_Enterprise_564

US_Startup_795

EMEA_Enterprise_270

US_Startup_318

EMEA_Startup_101

APAC_Enterprise_448

APAC_Enterprise_331

US_Startup_1147

EMEA_Startup_244

US_Startup_530

EMEA_Enterprise_92

US_Startup_379

APAC_Enterprise_56

US_Startup_218

APAC_Startup_10

US_Startup_921

EMEA_Enterprise_864

US_Enterprise_386

US_Enterprise_566

APAC_Startup_75

US_Startup_409

APAC_Enterprise_552

EMEA_Enterprise_51

US_Startup_125

US_Startup_399

APAC_Startup_95

APAC_Enterprise_58

US_Startup_200

APAC_Startup_43

US_Enterprise_695

APAC_Enterprise_106

APAC_Enterprise_682

APAC_Enterprise_597

US_Startup_352

EMEA_Enterprise_132

EMEA_Enterprise_50

US_Enterprise_556

US_Startup_972

EMEA_Enterprise_423

US_Enterprise_1300

Other

Word clouds usually look better with 20–100 words or phrases, but the format settings offer a wide range of flexibility. If you choose too many words, they can become too small to be legible, depending on the size of your display. By default, word clouds display 100 distinct words. To show more, change the format setting for **Number of words**.

Word cloud features

Working with items on sheets

Feature	Supported?	Comments	For more information
Changing the legend display	No		Legends on visual types in Quick
Changing the title display	Yes		Titles and subtitles on visual types in Quick
Changing the axis range	Not applicable		Range and scale on visual types in Quick
Changing the visual colors	Yes	To change the color, choose a word and then choose a color.	Colors in visual types in Quick
Focusing on or excluding elements	Yes		Focusing on visual elements Excluding visual elements
Sorting	Yes		Sorting visual data in Amazon Quick
Performing field aggregation	Yes	You can't apply aggregation to the field that you choose for Group by . You must apply an aggregation to the field that you choose for Size .	Changing field aggregation
Adding drill-downs	Yes	You can add drill-down levels to the Group by field well.	Adding drill-downs to visual data in Quick Sight
Using format options	Yes	You can choose to allow vertical words, emphasize scale, use a fluid layout, use lowercase, and set the amount of padding	Formatting in Amazon Quick

Feature	Supported?	Comments	For more information
		between words. You can set the maximum string length for the word cloud (default is 40). You can also choose the number of words for the Group by field (default is 100; maximum is 500).	
Showing totals	No		Formatting in Amazon Quick

Creating a word cloud

Use the following procedure to create a word cloud.

To create a word cloud

1. On the analysis page, choose **Visualize** on the tool bar.
2. Choose **Add** on the application bar, and then choose **Add visual**.
3. On the **Visual types** pane, choose the word cloud icon.
4. From the **Fields list** pane, drag the fields that you want to use to the appropriate field wells. Typically, you want to use dimension or measure fields as indicated by the target field well. If you choose to use a dimension field as a measure, the **Count** aggregate function is applied by default.

To create a word cloud, add a dimension to the **Group by** field well. Optionally, add a measure to the **Size** field well.

5. (Optional) Add drill-down layers by dragging one or more additional fields to the **Group by** field well. For more information about adding drill-downs, see [Adding drill-downs to visual data in Quick Sight](#).

Formatting in Amazon Quick

You choose from a variety of options to format and style your data visualizations. To format a visual, select the visual that you want to format and choose the **Format visual** icon on the upper-right corner of the visual. Once you have the format visual pane open, you can click on different

visuals and controls to view formatting data for the specific visual or control. For more information about formatting a visual control, see [Using a control with a parameter in Amazon Quick](#).

Use the following sections to format and style your content:

 Note

Any format changes applied from the field wells are applied only to the selected visual.

Topics

- [Analytics formatting per type in Quick](#)
- [Table and pivot table formatting options in Quick](#)
- [Adding data bars to tables in Quick](#)
- [Adding sparklines to tables in Quick](#)
- [Map and geospatial chart formatting options in Quick](#)
- [Axes and grid lines on visual types in Quick](#)
- [Colors in visual types in Quick](#)
- [Working with field level coloring in Amazon Quick](#)
- [Conditional formatting on visual types in Quick](#)
- [KPI options](#)
- [Labels on visual types in Quick](#)
- [Formatting visual numeric data based on language settings in Quick](#)
- [Legends on visual types in Quick](#)
- [Line and marker styling on line charts in Quick](#)
- [Missing data on visual types in Quick](#)
- [Reference lines on visuals types in Quick](#)
- [Formatting radar charts in Quick](#)
- [Range and scale on visual types in Quick](#)
- [Small multiples axis options](#)
- [Titles and subtitles on visual types in Quick](#)
- [Tooltips on visual types in Quick](#)

Analytics formatting per type in Quick

Use the following list to see what type of formatting works in a visualization during analysis:

- Bar charts (both horizontal and vertical) support the following formatting:
 - Customize, display, or hide title, field labels, and data labels
 - Customize, display, or hide legend (exception: simple charts without clustering or multiple measures don't show a legend)
 - Specify axis range and steps on the x-axis for horizontal bar charts, and on the y-axis for vertical bar charts
 - Choose how many data points to display on the x-axis for vertical bar charts, and on the y-axis for horizontal bar charts
 - Show or hide axis lines, axis labels, axis sort icons, and chart grid lines
 - Customize, display, or remove reference lines
 - Show or hide the "other" category

Horizontal bar charts support sorting on the y-axis and **Value**. Vertical bar charts support sorting on the x-axis and **Value**.

Stacked bar charts support showing totals.

- Box plots support the following formatting:
 - Customize, display, or hide title
 - Customize, display, or hide legend
 - Specify axis range and label tick on the x-axis and axis range and step on the y-axis
 - Show or hide axis lines, axis labels, axis sort icons, and chart grid lines
 - Choose how many data points to display on the y-axis.
 - Show or hide the "other" category
 - Add reference lines

Box plots support sorting on **Group by**.

- Combo charts support the following formatting:
 - Customize, display, or hide title, field labels, and data labels
 - Customize, display, or hide legend (exception: simple charts without clustering, stacking, or multiple measures don't show a legend)

- Specify axis range on bars and lines
- Synchronize the Y axes for both bars and lines into a single axis.
- Choose how many data points to display on the x-axis
- Show or hide axis lines, axis labels, axis sort icons, and chart grid lines
- Customize, display, or remove reference lines
- Show or hide the "other" category

Combo charts support sorting on the x-axis, **Bars**, and **Lines**.

- Donut charts support the following formatting:
 - Customize, display, or hide title, data labels, and legend
 - Customize, display, or hide the labels for group or color and value fields
 - Choose how many slices to display from **Group/Color**
 - Show or hide the "other" category

Donut charts support sorting on **Group/Color** and **Value**.

- Filled maps support the following formatting:
 - Customize, display, or hide title.
 - Customize, display, or hide the legend

Filled maps support sorting on **Location** and **Color**.

- Funnel charts support the following formatting:
 - Customize, display, or hide title, and data labels
 - Customize, display, or hide the labels for group or color and value fields
 - Choose how many stages to display in the **Group by** field
 - Show or hide the "other" category

Funnel charts support sorting on **Group by** and **Value**.

- Gauge charts support the following formatting:
 - Customize, display, or hide title. Display or hide axis labels.
 - Customize how to display the value or values: hidden, actual value, comparison
 - Choose the comparison method (available when you use two measures)
 - Choose the axis range and padding to display in the gauge chart

- Choose the arc style (degrees from 180 to 360) and arc thickness

Gauge charts don't support sorting.

- Geospatial charts (maps) support the following formatting:
 - Customize, display, or hide title and legend
 - Choose the base map image.
 - Choose to display map points with or without clustering.

Geospatial charts don't support sorting.

- Heat maps support the following formatting:
 - Customize, display, or hide title, legend, and labels
 - Choose how many rows and columns to display
 - Choose colors or gradients.
 - Show or hide the "other" category

Heat maps support sorting on **Values** and **Columns**.

- Histogram charts support the following formatting:
 - Customize, display, or hide title, field labels, and data labels
 - Specify axis range, scale, and steps on the y-axis
 - Choose how many data points to display on the x-axis
 - Show or hide axis lines, axis labels, axis sort icons, and chart grid lines

Histogram charts don't support sorting.

- Key performance indicators (KPIs) support the following formatting:
 - Customize, display, or hide title
 - Display or hide trend arrows and progress bar
 - Customize comparison method as auto, difference, percent (%), or difference as percent (%)
 - Customize primary value displayed to be comparison or actual
 - Conditional formatting

KPIs don't support sorting.

- Line charts support the following formatting:
 - Customize, display, or hide title, field labels, and data labels

• Customize, display, or hide legend (exception: simple charts don't show a legend)

- Specify axis range and steps (on y-axis)
- Choose how many data points to display on the x-axis
- Show or hide axis lines, axis labels, axis sort icons, and chart grid lines
- Customize, display, or remove reference lines
- Customize the styling of lines and the markers for data points on a line
- Show or hide the "other" category, except when the x-axis is a date

Line charts support sorting on the x-axis and **Value** for numeric purposes only.

- Pie charts support the following formatting:
 - Customize, display, or hide title, data labels, and legend
 - Customize, display, or hide the labels for group or color and value fields
 - Show metrics as values, percentages, or both
 - Choose how many slices to display from the **Group/Color** field
 - Show or hide the "other" category

Pie charts support sorting on **Value** and **Group/Color**.

- Pivot tables support the following formatting:
 - Customize, display, or hide title
 - Customize, display, or hide the labels for column, row, and value fields
 - Customize the font sizes for table headers and cells/body
 - Display or hide totals and subtotals on rows or columns
 - Custom labels for totals or subtotals
 - Choose additional styling options: fit table to view, hide +/- buttons, hide column field names, hide duplicate label when using single metric
 - Conditional formatting

Pivot tables support sorting on **Column** and **Row**. For more information on sorting pivot table data, see [Sorting pivot tables in Quick](#).

- Scatter plots support the following formatting:
 - Customize, display, or hide title, legend, field labels, and data labels
 - Customize, display, or remove reference lines

- Show or hide axis lines, axis labels, axis sort icons, and chart grid lines

Scatter plots don't support sorting.

- Tables support the following formatting:
 - Customize, display, or hide title, legend, and columns
 - Customize, display, or hide the column names for group-by and value fields
 - Customize the font sizes for table headers and cells/body
 - Display or hide totals at the top or bottom of the table
 - Provide a custom label for totals
 - Add conditional formatting

Tables support sorting on **Group by** and **Value**.

- Tree maps support the following formatting:
 - Customize, display, or hide title and legend
 - Customize, display, or hide the labels for group-by, size, and color fields
 - Choose colors or gradients.
 - Choose how many squares to display from the **Group by** field
 - Show or hide the "other" category

Line charts support sorting on **Size**, **Group by** and **Color**.

- Waterfall charts support the following formatting:
 - Customize, display, or hide title or subtitle
 - Customize the total label
 - Specify x-axis label size and orientation and y-axis label range and orientation.
 - Show or hide axis lines, axis labels, axis sort icons, and chart grid lines
 - Show or hide the "other" category
 - Customize the legend size and position.
 - Customize and display or hide data labels.

Waterfall charts support sorting on **Category** and **Value**.

- Word clouds support the following formatting:
 - Customize, display, or hide title
 - Customize the word color, and the number of words to display from the **Group by** field

- Show or hide the "other" category
- Choose additional styling options: allow vertical words, emphasize scale, or work with fluid layout, lowercase, padding level, or maximum string length

Word clouds support sorting on **Group by**.

Table and pivot table formatting options in Quick


You can customize tables and pivot tables in Quick to meet your business needs. You can customize table headers, cells, and totals by specifying the color, size, wrap, and alignment of text in each. You can also specify the height of rows in a table, add borders and grid lines, and add custom background colors. In addition, you can customize how to display totals and subtotals.

If you have applied conditional formatting to a table or pivot table, it takes precedence over any other styling you configure.

When you export table or pivot table visuals to Microsoft Excel, the formatting customizations that you applied to the visual aren't reflected in the downloaded Excel file.

To format a table or pivot table

- In your analysis, choose the table or pivot table that you want to customize, and then choose the **Format visual** icon.



The **Properties** pane opens at left.

Following, you can find descriptions for options for customizing each area of your table or pivot table in the **Properties** pane.

Topics

- [Headers](#)
- [Cell formatting](#)
- [Totals and subtotals](#)
- [Row and column size in tables and pivot tables in Quick](#)
- [Customize pivot table data](#)

Headers

Expand all headers

You can choose to expand all headers in a pivot table to show all child and grandchild rows of a header.

To expand all headers of a pivot table

1. On the visual that you want to change, select any header to open the **On-visual** menu.
2. Choose **Expand all below**.

Header height

You can customize table header height.

To customize the height of headers in a table

1. In the **Properties** pane, choose **Headers**.
2. For **Row height**, enter a number in pixels. You can enter a whole number from 8 through 500.

To customize the height of headers in a pivot table

1. In the **Properties** pane, choose **Headers**.
2. In the **Columns** section, for **Row height**, enter a number in pixels. You can enter a whole number from 8 through 500.

Header text

You can customize table header text.

To customize header text in a table

1. In the **Properties** pane, choose **Headers**.
2. Navigate to the **TEXT** section and do one or more of the following:
 - To change the color of the header text, choose the color swatch underneath **Text styling**, and then choose the color that you want the table text to be.

- To change the font or font size of the header text, open the **Font** or **Font size** dropdown and choose the font or font size that you want.
- To bold, italicize, or underline the header text, choose the appropriate icon from the style bar.
- To wrap text in headers that are too long to fit, select **Wrap text**. Wrapping text in a header doesn't automatically increase the height of the header. Follow the previous procedure for increasing header height.
- To change the horizontal alignment of text in the header, choose a horizontal alignment icon. You can choose left alignment, center alignment, right alignment, or automatic alignment.
- To change the vertical alignment of text in the header, choose a vertical alignment icon. You can choose top alignment, middle alignment, or bottom alignment.

To customize header text in a pivot table

1. In the **Properties** pane, choose **Headers**.

The Headers section expands to show options for customizing column and row headers.

2. In the **Headers** section, do one or more of the following:
 - To apply row styling to field names of the rows or columns, choose **Style rows label** or **Style columns label** depending on the label that you want to customize.
 - To customize the header font, navigate to the **TEXT** subsection of the **Rows** or **Columns** section and do one or more of the following:
 - To change the color of the header text, choose the color swatch underneath **Text styling**, and then choose the color that you want the table text to be.
 - To change the font or font size of the header text, open the **Font** or **Font size** dropdown and choose the font or font size that you want.
 - To bold, italicize, or underline the header text, choose the appropriate icon from the style bar.
 - To change the horizontal alignment of text in the header, choose an alignment icon. You can choose left alignment, center alignment, right alignment, or automatic alignment. You can choose a horizontal alignment for column headers in the **Columns** section, and for row headers in the **Rows** section.

- To change the vertical alignment of text in the header, choose an alignment icon. You can choose top alignment, middle alignment, or bottom alignment. You can choose a vertical alignment for column headers in the **Columns** section, and row headers in the **Rows** section.
- To hide the rows label or column field names, choose the eye icon next to **Rows label** or **Column field names**.

Header background color

You can customize table headers' background color.

To customize the background color of table headers

1. In the **Properties** pane, choose **Headers**.
2. For **Background**, choose the background color icon, and then choose a color. You can choose one of the provided colors, reset the header text color to the default color, or create a custom color.

To customize the background color of pivot table headers

1. In the **Properties** pane, choose **Headers**.

The **Headers** section expands to show options for customizing column and row headers.

2. In the **Columns** section, choose the background color icon, and then choose a color.
3. In the **Rows** section, choose the background color icon, and then choose a color.

Header borders

You can customize header borders' color.

To customize header borders in a table

1. In the **Properties** pane, choose **Headers**.
2. For **Borders**, do one or more of the following:
 - To customize the type of border that you want, choose a border type icon. You can choose no borders, horizontal borders only, vertical borders only, or all borders.
 - To customize the border thickness, choose a border thickness.

- To customize the border color, choose the border color icon, and then choose a color. You can choose one of the provided colors, reset the border color to the default color, or create a custom color.

To customize header borders in a pivot table

1. In the **Properties** pane, choose **Headers**.

The **Headers** section expands to show options for customizing column and row headers.

2. In the **Columns** and **Rows** sections, for **Borders**, do one or more of the following:
 - To customize the type of border that you want, choose a border type icon. You can choose no borders, horizontal borders only, vertical borders only, or all borders.
 - To customize the border thickness, choose a border thickness.
 - To customize the border color, choose the border color icon, and then choose a color. You can choose one of the provided colors, reset the border color to the default color, or create a custom color.

Header styling options for hierarchy pivot tables

You can hide or rename the **Rows** label of a hierarchy pivot table.

To make changes to the Rows label of a hierarchy pivot table

1. Select the hierarchy pivot table that you want to change and open the **Format visual** menu.
2. In the **Headers** section, you can perform the following tasks
 - Choose **Hide rows label** to hide the **Rows** label from your pivot table.
 - For **Rows label**, enter the label that you want displayed on the pivot table.

Cell formatting

Row height

You can customize table row height.

To customize the height of rows in a table or pivot table

1. In the **Properties** pane, choose **Cells**.

The **Cells** section expands to show options for customizing cells.

2. For **Row height**, enter a number in pixels. You can enter a whole number from 8 through 500.

Cell text

You can customize the formatting for cell text within a table.

To format the cell text in a table or pivot table

1. In the **Properties** pane, choose **Cells**.

The **Cells** section expands to show options for customizing cells.

2. For **Text**, do one or more of the following:
 - To change the color of the cell text, choose the color swatch underneath **Text styling**, and then choose the color that you want the table text to be.
 - To change the font or font size of the cell text, open the **Font** or **Font size** dropdown and choose the font or font size that you want.
 - To bold, italicize, or underline the cell text, choose the appropriate icon from the style bar.
 - To wrap text in headers that are too long to fit, select **Wrap text**. Wrapping text in cells doesn't automatically increase the row height. Follow the previous procedure for increasing row height.
 - To change the horizontal alignment of text in cells, choose a horizontal alignment icon. You can choose left alignment, center alignment, right alignment, or automatic alignment. Horizontal alignment can only be configured for the **Rows** fields of a hierarchy pivot table.
 - To change the vertical alignment of text in cells, choose a vertical alignment icon. You can choose top alignment, middle alignment, bottom alignment, or automatic. For tabular pivot tables, the value for **Automatic** is vertical. For hierarchy pivot tables, the value for **Automatic** is middle.

TEXT

Text styling

 Amazon ... 

B *i* U

Alignment

   Auto

☒ Wrap text

Cell background color

You can customize table cells' background color.

To customize the background color of cells in a table or pivot table

1. In the **Properties** pane, choose **Cells**.

The **Cells** section expands to show options for customizing cells.

2. For **Background**, do one or more of the following:
 - To alternate background colors between rows, select **Alternate row colors**. Clearing this option means that all cells have the same background color.
 - If you choose to alternate background colors between rows, choose a color for **Odd rows** and a color for **Even rows** by choosing the background color icon for each and selecting a color. You can choose one of the provided colors, reset the background color to the default color, or create a custom color.
 - If you choose not to alternate background colors between rows, choose the background color icon and select a color for all cells. You can choose one of the provided colors, reset the background color to the default color, or create a custom color.

Cell borders

You can customize table cells' borders.

To customize the borders for cells in a table or pivot table

1. In the **Properties** pane, choose **Cells**.

The **Cells** section expands to show options for customizing cells.

2. For **Borders**, do one or more of the following:

- To customize the type of border that you want, choose a border type icon. You can choose no borders, horizontal borders only, vertical borders only, or all borders.
- To customize the border thickness, choose a border thickness.
- To customize the border color, choose the border color icon, and then choose a color. You can choose one of the provided colors, reset the border color to the default color, or create a custom color.

Totals and subtotals

On tables and pivot tables, you can configure the display of totals or subtotals. Tables can display totals at the top or the bottom of the visual. Pivot tables can display totals and subtotals on rows and columns.

Add totals and subtotals to tables and pivot tables in Quick

You can add total columns to your table and pivot table visuals. You can also add subtotal columns to your pivot table visuals.

To display or hide totals and subtotals for a pivot table

1. To display totals, open the **Properties** pane and choose **Total**.
 - To show totals for rows, toggle the **ROWS** switch on. Totals are displayed on the bottom row of the visual. Choose **Pin totals** to keep the totals visible as you scroll through the table.
 - To show totals for columns, toggle the **COLUMNS** switch on. Totals are displayed on the last column of the visual.
2. To display totals, open the **Properties** pane and choose **Subtotal**.
 - To show subtotals for rows, toggle the **ROWS** switch on. Totals are displayed on the bottom row of the visual.
 - To show subtotals for columns, toggle the **COLUMNS** switch on.
 - For **Level**, choose one of the following:
 - Choose **Last** to only show the subtotal of the last field in the chart's hierarchy. This is the default option.

- Choose **All** to show subtotals for every field.
- Choose **Custom** to customize which fields show subtotals.

After you add row totals to your table or pivot table visual, you can also choose to position the totals at the top or bottom of the visual. You can also change the position of column totals in pivot tables.

To position row or column totals in a table or pivot table

1. In the **Properties** pane, choose **Total**.
2. (Optional) For **Rows**, choose **Show totals**.
3. (Optional) For **Columns**, choose **Show totals**.
4. (Optional) In the **Rows** menu, open the **Position** dropdown and choose the position that you want the totals to be displayed. Choose **Top** to position totals at the top of the table, or **Bottom** to position totals at the bottom of the table.
5. (Optional) In the **Columns** menu, open the **Position** dropdown and choose the position that you want the totals to be displayed. Choose **Left** to position totals at the left of the table, or **Right** to position totals at the right of the table.

You can't change the position of the subtotals of a pivot table visual. If your pivot table uses a hierarchy layout, the subtotal rows are positioned at the top of the table. Tabular pivot table subtotals are displayed at the bottom of the table.

Customize labels for totals and subtotals

You can rename the totals in table and pivot table visuals to provide better context for account readers. By default, the totals and subtotals appear without a label.

To rename totals in a table or pivot table visual

1. In the **Properties** pane, choose **Total** or **Subtotal**.
2. For **Label**, enter a word or short phrase that you want displayed for the total.

In pivot tables, you can also add labels to column totals and subtotals. To do so, enter a word or short phrase for **Label** in the **Columns** section.

3. (Optional) For tabular pivot tables, you can also add group names to subtotals. To add a group name to row subtotals, choose the **Plus (+)** icon next to the **Label** field to add the group name parameter that you want. You can also enter a word or short phrase to this field.

You can also make changes to the text size and font color of the total and subtotal labels of your table and pivot table visuals.

To format totals and subtotals text

1. In the **Properties** pane, choose **Total** or **Subtotal**.
2. For **Text**, do one or more of the following.
 - To change the color of the total or subtotal text, choose the color swatch underneath **Text styling**, and then choose the color that you want the table text to be.
 - To change the font or font size of the total or subtotal text, open the **Font** or **Font size** dropdown and choose the font or font size that you want.
 - To bold, italicize, or underline the total or subtotal text, choose the appropriate icon from the style bar.

In pivot tables, you can also add format text for column totals and subtotals. To do so, repeat the above steps in the **Columns** section.

Totals and subtotals background color

To customize the background color for totals and subtotals

1. In the **Properties** pane, choose **Total** or **Subtotal**.
2. For **Background**, choose the background color icon, and then choose a color. You can choose one of the provided colors, reset the background color to the default color, or create a custom color.

In pivot tables, you can also add background colors for column totals and subtotals. To do so, choose a the background color icon for **Background** in the **Columns** section.

Totals and subtotals borders

To customize the borders for totals and subtotals

1. In the **Properties** pane, choose **Total** or **Subtotal**.
2. For **Borders**, do one or more of the following:
 - To customize the type of border that you want, choose a border type icon. You can choose no borders, horizontal borders only, vertical borders only, or all borders.
 - To customize the border thickness, choose a border thickness.
 - To customize the border color, choose the border color icon, and then choose a color. You can choose one of the provided colors, reset the border color to the default color, or create a custom color.

In pivot tables, you can also add borders for column totals and subtotals. To do so, repeat the above steps in the **Columns** section.

Applying totals and subtotals styling to cells

In pivot tables, you can apply any text, background color, and border styling you apply to totals to cells in that same column or row. Row subtotals appear differently depending on the layout that your pivot table uses. For tabular pivot tables, explicit subtotal headers appear on the visual. For hierarchy pivot tables, explicit subtotal headers do not appear. Instead, authors apply subtotal styling to individual fields from the **Format visual** menu. Collapsed headers cannot be styled as subtotals.

To apply totals and subtotals styling to cells

1. In the **Properties** pane, choose **Total** or **Subtotal**.
2. For **Apply styling to**, choose the visual that you want to apply subtotal styling to. You can choose from the following options.
 - **None**– Removes styling options from all cells.
 - **Headers only**– Applies styling options to all headers in the pivot table.
 - **Cells only**– Applies styling options to all cells that aren't headers in the pivot table.
 - **Headers and cells**– Applies styling options to all cells in the pivot table.

Row and column size in tables and pivot tables in Quick

Authors and readers can resize rows and columns in a table or pivot table visual. They can adjust both row height and column width. Authors can also set the default column width for columns in a pivot table visual.

To resize a row in a table or pivot table

- In the table or pivot table visual, hover your cursor over the line that you want to resize until you see the horizontal cursor appear. When it appears, select the line and drag it to a new height.

You can adjust the row height by selecting the horizontal lines on cells and row headers.

		Region		
		EMEA	APJ	AMER
Industry	Product	Sales	Sales	Sales
Communications	Big Ol Database	5,178.58	287.91	3,617.66
	ChatBot Plugin	550.1	94.67	490.45
	ContactMatcher	14,814.21	1,984.05	10,844.09
	Data Smasher	7,911.58	1,864.1	3,418.41
	FinanceHub	7,811.35	1,203.52	7,866.2
	Marketing Suite	1,359.89	362.94	1,668.13
	Marketing Suite ...	4,816.37	2,561.22	3,795.39
	OneView	3,051.7	796.06	1,527.89
	SaaS Connector ...	2,769.74	1,328.2	1,785.16
	SaaS Connector ...	481.34	362.46	195.36
	Site Analytics	11,169.85	1,380.05	6,720.26
	Storage	98.31	25.25	102.18
	Support	4,414.38	1,265.26	987.94
	Alchemy	5,379.87	20,459.89	13,999.96
Consumer Products	Big Ol Database	836.08	868.8	2,966.2
	...	...	...	...

To resize a column width in a table or pivot table

- In the table or pivot table visual, hover your cursor over the line that you want to resize until you see the vertical cursor appear. When it appears, select the line and drag it to a new width.

You can adjust the column width by selecting the vertical lines on cells, column headers, and row headers.

Sum of Sales by Industry, Product, and Region

Industry	Product	Region		
		EMEA	APJ	AMER
		Sales	Sales	Sales
☐ Communications	Big Ol Database	5,178.58	287.91	3,617.66
	ChatBot Plugin	550.1	94.67	490.45
	ContactMatcher	14,814.21	1,984.05	10,844.09
	Data Smasher	7,911.58	1,864.1	3,418.41
	FinanceHub	7,811.35	1,203.52	7,866.2
	Marketing Suite	1,359.89	362.94	1,668.13
	Marketing Suite ...	4,816.37	2,561.22	3,795.39
	OneView	3,051.7	796.06	1,527.89
	SaaS Connector ...	2,769.74	1,328.2	1,785.16
	SaaS Connector ...	481.34	362.46	195.36
	Site Analytics	11,169.85	1,380.05	6,720.26
	Storage	98.31	25.25	102.18
	Support	4,414.38	1,265.26	987.94
☐ Consumer Products	Alchemy	5,379.87	20,459.89	13,999.96
	Big Ol Database	836.08	868.8	2,966.2

To set the default column width for columns in a pivot table

1. Select the pivot table that you want to change and open the **Format visual** menu.
2. In the **Pivot options** section, navigate to the **Value column width (pixels)** field and enter the default value that you want in pixels.

Customize pivot table data

You can customize how Quick readers view pivot tables so that they are easier to read and understand at a glance. You can choose to hide a pivot table's plus and minus icons, hide columns that only have a single-metric value, and hide collapsed columns from view. These options can help Quick authors remove clutter from their pivot tables and provide an easier reader experience for Quick users. This is not the same as choosing a pivot table layout. For more information on pivot table layout options, see [Choosing a layout](#).

These options can also be accessed from the **Combined row fields menu** of a pivot table. The layout that you choose for your pivot table determines how this menu is accessed. For more information on accessing the **Combined row fields** menu, see .

To make changes to a pivot table's layout

1. In the **Format visual** pane, choose **Pivot options**.
2. In the **Pivot options** menu, select the following options to customize the view:
 - **Hide +/- buttons** – Hide the plus and minus icons from your pivot table by default. Readers can still choose to show the plus and minus icons and expand or collapse columns and rows.
 - **Hide single metric** – Hide columns that only have a single metric value.
 - **Hide collapsed columns** – Automatically hide all collapsed columns in a pivot table. This option is only available for tabular pivot tables.

Adding data bars to tables in Quick

You can use data bars to add visual context to your table visuals in Amazon Quick. By injecting color into your tables, data bars can make it easier to visualize and compare data in a range of fields. *Data bars* are bars of different colors or shades that you add to the cells of a table. The bars are measured relative to the range of all cells in a single column, which is similar to a bar chart. You can use data bars to highlight a fluctuating trend, such as profit per quarter during the year.

You can only apply data bars to fields that are added to the **Values** field well of the visual. You can't apply data bars to items that are added to group bys.

You can create up to 200 different data bar configurations for a single table.

Sum of Sales (Sum) and Sum of Profit (Sum) by Industry

industry	sales (SUM)	profit (SUM)
Agriculture	166,710.24	-34,659.2
Communications	1,177,239.36	110,953.52
Consumer Products	1,738,051.36	411,494.66
Energy	2,399,247.52	344,129.9
Finance	4,193,206.24	508,255.05
Healthcare	1,904,002.08	195,355.16
Industrial	253,124.96	55,607.64
Misc	1,891,786.24	222,764.91
Other	115,406.88	7,703.52
Retail	3,089,122.08	452,895.46
Tech	2,781,517.28	347,750.8
Transportation	1,691,424.16	188,505.49

To add data bars to a table

1. On the analysis page, choose the visual that you want to format.

2. On the menu in the upper-right corner of the visual, select the **Format visual** icon. The **Format visual** pane opens.
3. In the **Properties** pane, open the **Visuals** dropdown list and choose **ADD DATA BARS**.
4. In the **Data bars** popup that appears, choose the value field that you want represented by the data bars. You can only choose from fields that are added to the **Values** field well of the visual.
5. (Optional) Choose the icon labeled **Positive color** to select the color that you want to represent positive value data bars. The default color is green.
6. (Optional) Choose the icon labeled **Negative color** to select the color that you want to represent negative value data bars. The default color is red.

When you create data bars, they are named for the field values that they are representing. For example, if you add data bars to represent the profit of a product over time, the data bar configuration is labeled "Profit". In the **Visuals** pane of the **Properties** menu, data bars are listed in the order that they are created.

To remove data bars from a visual

1. On the menu in the upper-right corner of the visual, select the **Format visual** icon. The **Properties** pane opens.
2. In the **Properties** pane, open the **Visuals** dropdown list and choose the data bar that you want to remove.
3. Choose **REMOVE DATA BARS**.

Adding sparklines to tables in Quick

Sparklines are small inline charts that display trends directly within table cells, helping readers quickly identify patterns and seasonality without leaving the table view. Use sparklines when you need compact trend visualization alongside your tabular data.

To apply sparklines to a table

1. On the analysis page, choose the table visual that you want to format.
2. On the menu in the upper-right corner of the visual, select the **Format visual** icon. The **Format visual** pane opens.
3. In the **Properties** pane, open the **Visuals** dropdown list and choose **APPLY SPARKLINES**.
4. In the sparklines editing pane, configure the data settings:

- For **Value column**, choose the measure field that you want the sparkline to represent. Fields already used by another sparkline or data bar are not available.
 - For **X-axis field**, choose the dimension field to plot along the horizontal axis. The X-axis field must not be the same as a field in the **Group by** field well. You can also configure the sort direction and time granularity (for date/time fields) of the X-axis field.
5. (Optional) Expand the **Presentation** section to customize the sparkline appearance. See [Sparkline options](#) for details.
6. (Optional) Configure marker visibility. All markers are hidden by default. You can choose to show:
- **All points** – Show a marker on every data point.
 - **Max value** – Show a marker on the highest value.
 - **Min value** – Show a marker on the lowest value.
7. Choose **Apply**.

The sparkline is named after the value field it represents (for example, "Profit"). Sparklines appear in the **Visuals** pane in the order they are created.

Sparkline options

The following table describes the sparkline presentation options.

Setting	Options	Default	Description
Y-axis behavior	Shared, Independent	Shared	Shared uses the same Y-axis scale across all rows for easy comparison. Independent scales each row separately to highlight individual trend shapes.
Visual type	Line, Area line	Line	Area line adds a shaded area beneath the line.
Line color	Color picker	Theme color	Custom color for the sparkline line.

Setting	Options	Default	Description
Line interpolation	Linear, Smooth, Stepped	Linear	Controls how points are connected.

Editing and removing sparklines

To edit a sparkline, open the **Visuals** dropdown in the **Format visual** pane and choose the edit icon next to the sparkline you want to modify. Update the settings and choose **Apply**.

To remove a sparkline, open the edit pane for the sparkline and choose **Delete**.

Automatic removal

Quick automatically removes sparklines when field changes make them invalid:

- All **Group by** fields are removed – all sparklines are removed.
- A sparkline's value column is removed from the **Values** field well – that sparkline is removed.
- A sparkline's X-axis field is added to the **Group by** field well – that sparkline is removed.

A notification appears when a sparkline is automatically removed.

Sparkline limitations


Consider the following when working with sparklines:

- **Maximum sparklines per table** – Up to 3 sparkline columns per table visual
- **Maximum data points** – Up to 52 data points per sparkline. If your data exceeds this limit, Quick displays the last 52 data points according to your X-axis sort order.
- **Field requirements** – At least one field in the **Group by** field well and one field in the **Values** field well
- **X-axis constraint** – The X-axis field cannot be the same as any **Group by** field
- **Exclusive value column usage** – A value column cannot be used by both a sparkline and a data bar
- **Export support** – Sparklines are included in PDF exports but not in CSV or Excel exports
- **Filter behavior** – Filters applied to the table also filter sparkline data

Map and geospatial chart formatting options in Quick

In Amazon Quick, you can choose from multiple formatting options for your maps and geospatial charts. You can view formatting options by opening the **Properties** pane from the on-visual menu located at the top right of the currently selected geospatial map.

Quick authors and readers can also toggle the different formatting options of a geospatial map visual from the on visual menu.



Topics

- [Base maps on geospatial maps in Quick](#)
- [Geospatial heatmaps in Amazon Quick](#)
- [Marker clustering on geospatial point maps in Quick](#)

Base maps on geospatial maps in Quick


When you create a map visual in Quick, you can change the base of the map. A *base map* is the style of map that appears beneath your data on a map. An example is a satellite view versus a street view.

In Quick, there are four options for base maps: light gray canvas, dark gray canvas, streets, and imagery. The following list contains an example of each base map option:

⚠ Important

Only the light gray canvas is supported in the Asia Pacific (Mumbai) AWS Region (ap-south-1).


- Light gray canvas




- Dark gray canvas



- Streets



- Imagery



Changing base maps

Use the following procedure to change a base map.

To change a base map


1. Create a point or filled map in an analysis. For more information, see [Creating maps and geospatial charts](#).
2. On the map visual, choose the **Format visual** icon.
3. In the **Properties** pane that opens, choose the **Base map** section and then choose the base map that you want.

Geospatial heatmaps in Amazon Quick

Use geospatial heatmaps to reveal patterns of marker concentration in your geospatial visuals. Heat maps display concentrations of data points using a colored overlay that highlights the intensity or concentration of the visual's markers.

Field wells

Sheet 1 ▾ +




To turn a geospatial map into a heat map

1. Open your analysis and choose the geospatial map that you want to format. When you select a visual, it displays with a highlight around it.
2. To open the formatting pane, select the **Format visual** icon from the on-visual menu.
3. On the formatting pane at left, choose **Points**.
4. Choose **Heatmap**.
5. (Optional) For **Heatmap gradient**, choose a color that you want for the **High density** and **Low density** values.

Marker clustering on geospatial point maps in Quick

Use marker clustering to improve readability of collocated points on a map. Geospatial locations on point maps are represented using markers. Usually, there is one marker per data point. However, if there are too many markers close together, the map becomes difficult to read. To make it easier to interpret the map, you can enable marker clustering to represent groupings of locations on the

map. As the reader zooms in on the map, the clustered markers leave the area marker to display separately.



To add cluster points to a map

1. Open your analysis, and choose the geospatial map that you want to format. When you select a visual, it displays with a highlight around it.
2. To open the formatting pane, select the **Format visual** icon from the on-visual menu.
3. On the formatting pane at left, choose **Points**.
4. Choose one of the following options:
 - **Basic** – use the default display setting for map points.
 - **Cluster points** – cluster map points together when there are many in one area.

Axes and grid lines on visual types in Quick

When you create a chart in Quick, axis lines, axis labels, axis sort icons, and grid lines are added to the chart automatically. You can format your visuals to show or hide these if you want, and also customize the axis label size and orientation.

You can format axis lines, grid lines, and axis labels and axis sort icons for the following chart types:

- Bar charts

- Box plot charts
- Combo charts
- Histograms
- Line charts
- Scatter plots
- Waterfall charts

To format axis lines, axis labels, and grid lines in a chart

1. On the analysis page, choose the visual that you want to format.
2. On the menu in the upper-right corner of the visual, select the format visual icon.

The **Properties** pane opens at left.

To show or hide axis lines

1. In the **Properties** pane, choose the axis that you want to format.
2. Choose **Show axis line**. Clear the check box to hide the axis line for the chosen axis. Select the check box to show it.

To customize axis titles

1. In the **Properties** pane, choose the axis that you want to format.
2. Choose **Show title**. Clear the check box to hide the axis title and drop-down caret icon for the chosen axis. Select the check box to show them.
3. To change the title from the default field name, enter a title in the text box.

Note

In addition to the chart types listed previously in this topic, you can also customize the axis titles in pie charts, donut charts, funnel charts, heat maps, and tree maps.

To modify axis font settings

1. In the **Properties** pane, choose the axis that you want to format.
2. Adjust the following properties:
 - **Font family**
 - **Text size**
 - **Style** (Bold, Italic, Underline)
 - **Color**

Note

- **Underline** is supported for axis titles, but not for axis labels
- Different chart types use different terminology:
 - **Bar/Line charts** - **X-axis** and **Y-axis**
 - **Pie charts** - **Values**
 - **Heat maps** - **Rows** and **Columns**

To show or hide the sort icon

1. In the **Properties** pane, choose the axis that you want to format.
2. Choose **Show sort**. Clear the check box to hide the sort icon for the chosen axis. Select the check box to show it.

When you choose to remove the sort icon, the sort icon is removed from the axis. Any sorts that were applied to the visual before removing the icon are not removed from the visual.

Note

In addition to the chart types listed previously in this topic, you can also show or hide the sort icon in pie charts, donut charts, funnel charts, heat maps, and tree maps.

To show or hide the data zoom

1. In the **Properties** pane, choose **X-axis**.
2. Choose **Show data zoom**. Clear the check box to hide the data zoom. Select the check box to show it.

The data zoom bar appears automatically on charts with an X-axis that contain more than one data point. Adjust the bar from the left and right to zoom to specific data points in the chart.

Note

If you zoom in or out using the data zoom bar, and then choose to hide the data zoom bar, the zoom position isn't maintained. The visual zooms completely out to include all data points. Showing the data zoom again returns the visual to its previous state.

To show or hide axis labels

1. In the **Properties** pane, choose the axis that you want to format.
2. Choose **Show labels**. Clear the check box to hide the axis labels for the chosen axis. Select the check box to show it.

To change the label size

1. In the **Properties** pane, choose the axis that you want to format.
2. For **Label size**, choose a size.

To change the label orientation

1. In the **Properties** pane, choose the axis that you want to format.
2. For **Label orientation**, choose an orientation.

To show or hide grid lines

1. In the **Properties** pane, choose the axis that you want to format.
2. Choose **Show grid lines**. Clear the check box to hide grid lines for the chosen axis. Select the check box to show it.

Colors in visual types in Quick

You can change the color of one, some, or all elements on the following types of charts:

- Bar charts
- Donut charts
- Gauge charts
- Heat maps
- Line charts
- Scatter plots
- Tree maps


To change colors on bar charts, donut charts, gauge charts, line charts, and scatter plots, see [Changing colors on charts](#).

To change colors on heat maps and tree maps, see [Changing colors on heat maps and tree maps](#).


Changing colors on charts

You can change the chart color used by all elements on the chart, and also change the color of individual elements. When you set the color for an individual element, it overrides the chart color.


For example, suppose that you set the chart color to green.




All of the bars turn green. Even though you choose the first bar, the chart color applies to all the bars. Then you set the color for the **SMB** bar to blue.




Looking at the result, you decide that you need more contrast between the green and blue bars, so you change the chart color to orange. If you are changing the chart color, it doesn't matter which bar you choose to open the context menu from.



The **SMB** bar remains blue. This is because it was directly configured. The remaining bars turn orange.




When you change the color of an element that is grouped, the color for that element is changed in all of the groups. An example is a bar in a clustered bar chart. In the following example, Customer Segment is moved out of the **Y-axis** and into the **Group/Color** field well. Customer Region is added as the **Y-axis**. The chart color stays orange, and SMB stays blue for all Customer Regions.



If your visual has a legend that shows categories (dimensions), you can click on the values in the legend to see a menu of available actions. For example, suppose that your bar chart has a field in the **Color** or **Group/Color** field well. The bar chart menu displays the actions that you can choose by clicking or right-clicking on a bar, such as the following:

- Focusing on, or excluding, visual elements
- Changing colors of visual elements
- Drilling down into a hierarchy
- Custom actions activated from the menu, including filtering or URL actions

Following is an example of using the legend to change the color for a dimension.



Setting new colors for a visual

Use the following procedure to change the colors for a visual.

To change the colors for a visual

1. On the analysis page, choose the visual that you want to modify.
2. To change the chart color, choose any element on the visual, and then choose **Chart Color**.

To select elements, do the following:

- On a bar chart, choose any bar.
 - On a line chart, choose the end of a line.
 - On a scatter plot, choose an element. The field must be in the **Group/Color** section of **Field wells**.
3. Choose the color that you want to use. You can choose a color from the existing palette, or you can choose a custom color. To use a custom color, enter the hexadecimal code for that color.

All elements on the visual are changed to use this color, except for any that have previously had their color individually set. In that case, the element color overrides the chart color.

4. To change the color for a single element on the visual, choose that element, choose **Color <field name>**, and then choose the color that you want to use. You can choose a color from the existing palette, or you can choose a custom color. To use a custom color, enter the hexadecimal code for that color.

Repeat this step until you have set the color on all elements that you want to modify. To change the color back to the color it was originally, choose **Reset to default**.

Setting visual colors back to defaults

Use the following procedure to return to using the default colors on a visual.

To return to default colors on a visual

1. On the analysis page, choose the visual that you want to modify.
2. Choose **Chart Color**, choose any element on the visual, and then choose **Reset to Default**. Doing this changes the chart color back to the default color for that visual type.

All elements on the visual are changed to the default color for the visual type, except for any that have previously had their color individually set. In that case, the element color setting overrides the chart color setting.

3. To change the color for a single element back to the default, choose that element, choose **Color <field name>**, and then choose **Reset to Default**.

The default color for individual elements is the chart color if you have specified one, or the default color for the visual type otherwise.

Changing colors on heat maps and tree maps

To change the colors that display on a heat map or a tree map

1. Choose the heat map or tree map that you want to edit.
2. Choose **Expand** for the settings menu, and choose the cog icon to open the **Properties** panel.
3. For **Color**, choose the settings that you want to use:
4. For **Gradient color** or **Discrete color**, choose the color square next to the color bar, and then choose the color that you want to use. Repeat for each color square. The bar holds two colors by default.
5. Select the **Enable 3 colors** check box if you want to add a third color. A new square appears in the middle of the color bar.


You can enter a number that defines the midpoint between the two main gradient colors. If you add a value, the middle color represents the number you entered. If you leave this blank, the middle color acts like the other colors in the gradient.

6. Select the **Enable steps** check box if you want to limit the chart to the colors that you chose. Doing this changes the label on the color bar from **Gradient color** to **Discrete color**.
7. For **Color for Null Value**, choose a color to depict NULL values. This option is only available on heat maps.

Working with field level coloring in Amazon Quick

With field level coloring, you can assign specific colors to specific field values across all visuals in a Quick analysis or dashboard. Colors are assigned on a per-field basis to simplify the process of setting colors and ensure consistency across all visuals that use the same field. For example, let's say you're a shipping company that wants to create a set of visuals that track shipping rates in different regions. With field level coloring, you can assign each region a different color to represent the field across all visuals in an analysis or dashboard. This way, account readers quickly learn what field colors they're looking for and have an easier time finding the information that they need.

Quick authors can configure up to 50 field based colors per field. Colors that are defined at the visual level take precedence over field based colors. This means that if the author sets a color for a value on the visual, that color will override the field based colors configuration for that individual visual.



To apply field level coloring to a legacy account

1. In the **Fields** pane of the analysis, choose the ellipsis (three dots) next to the field that you want to assign a color to, and then choose **Edit field colors**.
2. In the **Edit field colors** pane that appears, choose the value that you want to assign a color to and choose the color that you want. You can apply colors to every value that appears in the **Field values** pane.
3. When you are finished assigning colors to the fields that you want, choose **Apply**.

If you want to reset the color value of a field, open the **Edit field colors** pane and choose the refresh icon next to the field that you want to reset. You can reset all color values in an analysis by choosing **RESET COLORS**.

You can view a list of unused colors that can be configured to new fields by choosing **Show unused colors** in the **Edit field colors** pane. When you reset a field's color, the discarded color is added to the **Unused colors** list and can be assigned to a new field.

Conditional formatting on visual types in Quick

In some visual types, you can add conditional formatting to highlight some of your data. The conditional formatting options currently supported include changing text or background color and using symbolic icons. You can use icons from the provided set, or you can use Unicode icons instead.

Conditional formatting is available on the following visuals:

- Gauge charts
- Key performance indicators (KPIs)
- Pivot tables
- Tables

For tables and pivot tables, you can set multiple conditions for fields or supported aggregations, along with format options to apply to a target cell. For KPIs and gauge charts, you can format the primary value based on conditions that are applied to any dimension in the dataset. For gauge charts, you can also format the foreground color of the arc based on conditions.

To use conditional formatting on a visual

1. On the analysis page, choose the visual that you want to format.
2. On the visual, open the context menu on the down icon at the upper-right. Then choose **Conditional formatting**.

Options for formatting display on the left. Choose one of the following:

- **For pivot tables** – Begin by choosing a measure that you want to use. You can set conditional formatting on one or more fields. The selection is limited to the measures that are in the **Values** field well.
 - **For tables** – Begin by choosing a field that you want to use. You can set conditional formatting on one or more fields. You can also choose to apply formatting to the entire row. Formatting the entire row adds an option to **Apply on top**, which applies the row formatting in addition to formatting added by other conditions.
 - **For KPIs** – Apply formatting to the primary value or the progress bar or both.
3. For the remaining steps in this procedure, choose the features that you want to use. Not all options are available for all visuals.
 4. (Optional) Choose **Add background color** to set a background color. If a background color is already added, choose **Background**.
 - **Fill type** – The background color can be **Solid** or **Gradient**. If you choose to use a gradient, additional color options display, enabling you to choose a minimum and maximum value for the gradient scale. The minimum value defaults to the lowest value, and the maximum value defaults to the highest value.

- **Format field based on** – The field to use when applying the format.
 - **Aggregation** – The aggregation to use (displays only the available aggregations).
 - **Condition** – The comparison operator to use, for example "greater than".
 - **Value** – The value to use.
 - **Color** – The color to use.
 - **Additional options:** In pivot tables, you can set what you want to format by choosing options from the context menu (...): **Values**, **Subtotals**, and **Totals**.
5. (Optional) Choose **Add text color** to set a text color. If a text color is already added, choose **Text**.
- **Format field based on** – The field or item to use when applying the format.
 - **Aggregation** – The aggregation to use (displays only the available aggregations). This option applies to tables and pivot tables.
 - **Condition** – The comparison operator to use, for example "greater than".
 - **Value** – The value to use.
 - **Color** – The color to use.
 - **Additional options:** In tables and pivot tables, you can set what you want to format by choosing options from the context menu (...): **Values**, **Subtotals**, and **Totals**.
6. (Optional) Choose **Add icons** to set an icon or icon set. If an icon is already added, choose **Icon**.
- **Format field based on** – The field or item to use when applying the format.
 - **Aggregation** – The aggregation to use (displays only the available aggregations). This option applies to tables and pivot tables.
 - **Icon set** – The icon set to apply to field in **Format field based on**. This option applies to tables and pivot tables.
 - **Reverse colors** – Reverses the colors of the icons for tables and pivot tables.
 - **Custom conditions** – Provides more icon options for tables and pivot tables.
 - **Condition** – The comparison operator to use.
 - **Value** – The value to use.
 - **Icon** – The icon to use. To choose an icon set, use the **Icon** symbol to choose the icons to use. Choose from the provided icon sets. In some cases, you can add your own. To use your own icon, choose **Use custom Unicode icon**. Paste in the Unicode glyph that you want to use as an icon. Choose **Apply** to save or choose **Cancel** to exit icon setup.

- **Color** – The color to use.
 - **Show icon only** – Replaces the value with the icon for tables and pivot tables.
 - **Additional options:**
 - In tables and pivot tables, you can set what you want to format by choosing options from the context menu (...): **Values**, **Subtotals**, and **Totals**.
 - In pivot tables, enabling **Custom conditions** activates preset conditional formatting that you can keep, add to, or overwrite with your own settings.
7. (Optional) Choose **Add foreground color** to set the foreground color of a KPI progress bar. If a foreground color is already added, choose **Foreground**.
- **Format field based on** – The field to use when applying the format.
 - **Condition** – The comparison operator to use.
 - **Value** – The value to use.
 - **Color** – The color to use.
8. When you are finished configuring conditional formatting, choose one or more of the following:
- To save your work, choose **Apply**.
 - To cancel selections and return to the previous panel, choose **Cancel**.
 - To close the settings panel, choose **Close**.
 - To reset all settings on this panel, choose **Clear**.

KPI options

You can customize KPIs in Amazon Quick to meet your business needs. You can add contextual sparklines or progress bars, assign primary and secondary values, and add conditional formatting to your KPIs.

To format a KPI in Quick, navigate to the KPI that you want to change and choose the **Format visual** icon to open the **Format visual**.

Use the following procedures to perform formatting tasks for KPIs.

Add a visual to a KPI

You can choose to add an area sparkline, a sparkline, or a progress bar to any KPI in Quick. Adding visuals to KPIs provides visual context to readers who are viewing KPI data. Use the following procedure to add a visual to a KPI.

To add a visual to a KPI

1. Navigate to the KPI that you want to change and open the format visual menu.
2. In the **Properties** menu, choose the **Visual** box to display a visual on your KPI chart.
3. (Optional) Open the **Visual** dropdown and choose the type of visual that you want to display on your KPI. You can choose to display an area sparkline, a sparkline, or a progress bar. To display a sparkline, make sure that your KPI has a value in the **Trend** field well. **Area sparkline** is the default value.
4. (Optional) To change the color of the sparkline, choose the color icon to the left of the **Visual** dropdown and choose the color that you want. Color formatting isn't supported for the progress bar.
5. (Optional) Choose **Add tooltip** to add a tooltip to the KPI visual.

Customizing primary and secondary values

Use the **Format visual** menu to customize the font, color, and to choose which primary value is displayed. You can also choose to display a secondary value.

To customize the primary and secondary values of a KPI

1. Navigate to the KPI that you want to change, open the **Format visual** menu, and navigate to the **KPI** section.
2. For **Primary value**, use the **Font** dropdown to choose the font size that you want. The default value is **Auto**.
3. (Optional) To change the color of the primary value's font, choose the color icon next to the **Font** dropdown, and then choose the color that you want.
4. For **Primary value displayed**, you can choose to display the actual value or the comparison value of the primary value.
5. To add a secondary value, choose **Secondary value**.

- a. (Optional) Use the **Font** dropdown to choose the font size that you want. The default value is **Extra large**.
- b. (Optional) To change the color of the secondary value's font, choose the color icon next to the **Font** dropdown, and then choose the color that you want.

Conditional formatting options for KPIs

Conditional formatting for KPIs is automatically set for comparison values. By default, positive values are represented in green and negative values are represented in red. You can customize the color values of these color values from the **Properties** pane.

To change the color of positive and negative values

1. In the **Properties** pane, open the **Conditional formatting** section and choose the comparison value that you want to change.
2. To change the color of the positive value, navigate to **Condition #1**, choose the **Color** icon, and then choose the color tht you want.
3. To change the color of the negative value, navigate to **Condition #2**, choose the **Color** icon, and then choose the color tht you want.
4. When you are finished making the changes that you want, choose **Apply**.

You can also add text colors and icons for the **Actual value** in thee **Conditional formatting** menu. To add a text color or icon to the actual value, choose **Add text color** or **Add icon** to set the new values.

Labels on visual types in Quick

Use the following procedure to customize, display, or hide the labels for a visual.

To customize, display, or hide the labels for a visual

1. On the analysis page, choose the visual that you want to format. You can change the labels by choosing the label directly on the visual, and choosing **Rename**. To revert to the default name, delete your entry.
2. To see more options, choose the on-visual menu from the down icon at the upper-right corner of the visual, and then choose the **Format visual** icon.

For pivot tables, you can relabel row names, column names, and value names. Additionally, under **Styling**, you can choose to hide columns labels or metric labels (for single metrics only).

You can add the same value to the same visual multiple times. You can do so to show the same value with different aggregations or table calculations applied. By default, the fields all display the same label. You can edit the names by using the **Properties** panel, which you open by choosing the **V**-shaped icon at top right.

3. On the **Properties** pane, enable or disable **Show title**. This option removes the axis title.
4. Close the **Properties** pane by choosing the **X** icon in the upper-right corner of the pane.

Data labels on visual types in Quick

To customize data labels on a visual, you can use the **Properties** pane to show data labels, and then use the settings to configure them. Data label customization is supported on bar, line, combo, area, scatterplot, donut, boxplot, waterfall, heatmap, treemap, histogram, funnel, sankey, gauge, radar, and pie charts.

You can customize the following options:

- Position, which determines where the label appears in relation to the data point (for bar, combo, and line charts):
 - For vertical bar charts, you can customize to set position:
 - Above bars
 - Inside of bars
 - Bottom of bars
 - Top of bars
 - For horizontal bar charts, you can customize to set position:
 - Right of bars
 - Inside of bars
 - For line charts, you can customize to set position:
 - Above lines
 - Left or right of points on lines
 - Below lines
 - For scatter charts, you can customize to set position:

- Above points
- Left or right of points
- Below points
- Font size and color (for bar, combo, line, scatter, and pie charts)
- Label pattern, which determines how data is labeled (for bar, combo, line, and scatter charts):
 - For bar, combo, and scatter charts, you can label:
 - All
 - By group or color
 - For line charts, the following label options are available:
 - All
 - By group or color
 - Line ends
 - Minimum or maximum value only
 - Minimum and maximum values
 - For pie charts, the following label options are available:
 - Show category
 - Show metric
 - Choose to show the metric label as value, percent, or both
- Group selection (for bars and lines, when the label pattern is "by group/color")
- Allow labels to overlap (for bars and lines), for use with fewer data points
- For vertical bar, combo, and line charts, labels that are too long are angled by default. You can configure the degree of angle under the **X-axis** settings.

Note

If you add more than one measure to an axis, the data label displays the formatting for the first measure only.

To configure data labels

1. On the analysis page, choose the visual that you want to format.

2. Choose the on-visual menu from the down icon at the upper-right corner of the visual, and then choose the **Format visual** icon.
3. On the **Properties** pane, choose **Data Labels**.
4. Enable **Show data labels** to show and customize labels. Disable this option to hide data labels.
5. Choose the settings that you want to use. The settings offered are slightly different for each chart type. To see all available options, see the list before this procedure.

You can immediately view the effect of each change on the visual.

6. To modify the data label font settings, adjust the following properties:
 - **Font family**
 - **Text size**
 - **Style** (Bold, Italic)
 - **Color**
7. Close the **Properties** pane by choosing the X icon in the upper-right corner of the pane.

Formatting visual numeric data based on language settings in Quick

In Amazon Quick, you can choose how your numeric data values appear in visuals so that they align with the regional language that you have chosen.

As a Quick author, you can choose the language formatting that best fits your audience. Amazon Quick configures numeric data languages at the analysis level based on the language that you have chosen to view Quick in. You can change the format of numbers, currencies, and dates. You can change your Quick language settings in the **Language** dropdown list of the Quick **User** menu in the top-right corner. You can change the language formatting for a field across every visual in a sheet, or you can change the language formatting at the individual visual level.


To change the numeric language formatting of all visuals in an analysis

1. On the **Visuals** pane of the analysis that you want to change, choose the more actions (three dots) icon next to the field that you want to change. From the menu that appears, open the **Format** dropdown list, and then choose **More formatting options**.
2. In the **Format data** pane that appears on the left, choose **Apply language format**.


You can reset the default language format of the field by reopening the **Format data** menu and choosing **Reset to defaults**. The default language format is American English.

To change the numeric language formatting of a single visual in an analysis

1. On the analysis page, choose the visual that you want to modify.
2. Navigate to the **Format data** pane using one of the following options:
 - On the visual that contains the data that you want to change, select the field that you want to change, open the **Format** dropdown list, and then choose **More formatting options**.



- In the **Field wells** section of the analysis, open the dropdown next to the field that you want to change. Open the **Format** menu, and choose **More formatting options**.



3. In the **Format data** pane that appears, choose **Apply language format**.

You can reset the default language format of the visual by reopening the **Format data** menu and choosing **Reset to defaults**. The default language format is American English.

Legends on visual types in Quick

The *visual legend* helps you identify what a visual element represents by mapping its value to a color. By default, the visual legend displays to the right of the visual. You can choose to hide or display the visual legend, and format the legend title and position. You can also customize the font settings for the legend title and items.

To display or hide a visual legend

1. Sign in to Quick at <https://quicksight.aws.amazon.com/>.
2. On the analysis page, choose the visual that you want to format.
3. Choose the visual that you want to format, and then choose the **Properties** icon to open the Properties pane.
4. Toggle the **Legend** on to display the visual's legend. When shown, the legend displays the values in alphabetical order. To hide the legend, toggle the **Legend** switch off.

To customize a visual legend

1. Open the Properties pane and expand the **Legend** section.
2. Use the **Position** dropdown to customize the position of the legend in the visual.
3. For **Legend title**, enter a custom name for the legend and perform all or some of the following actions:
 - a. (Optional) To change the color of the legend title, choose the color swatch underneath the legend title, and then choose the color that you want the legend title to be.
 - b. (Optional) To change the font or font size of the legend title, open the **Font** or **Font size** dropdown and choose the font or font size that you want.
 - c. (Optional) To bold, italicize, or underline the legend title, choose the appropriate icon from the style bar.
4. For **Legend item**, perform all or some of the following actions:
 - a. (Optional) To change the color of the legend item font, choose the color swatch, and then choose the color that you want the legend title to be.
 - b. (Optional) To change the font or font size of the legend item, open the **Font** or **Font size** dropdown and choose the font or font size that you want.

- c. (Optional) To bold, italicize, or underline the legend item font, choose the appropriate icon from the style bar.
5. Choose the **X** icon at upper right to close the **Properties** pane.

Line and marker styling on line charts in Quick

In Quick line charts, you have multiple options to emphasize what you want readers to focus on: color, line style, and markers. You can use these options together or separately to help readers understand your line charts more quickly under different circumstances. For example, if some of your readers won't see color differences—perhaps because of color blindness or because of monochrome printing—you can use line patterns to distinguish one or more lines in a chart.

In other cases, you could use step lines to call attention to abrupt changes or intervals between changes in data. For example, let's say you build a chart showing the changing price of postage stamps in the US, and you want to emphasize the amount of increase in price over time. You can use a step line, which remains flat between data points until the next price change occurs. The data story about abrupt increases in price is more clear to the reader with a step line. If you wanted to show a story of gradual change over time, you'd be more likely to style the line with a smooth slope instead.

To customize the styling for a visualization

1. Open your analysis, and choose the chart that you want to format.
2. On the top right of the visual you want to format, select **Format visual**, which is represented by a pencil icon.
3. At left, choose **Data series**.
4. Choose one of the following options:
 - **Base style** – to edit the styling of all lines and markers on the chart
 - **Select series to style** – to edit the styling of the field that you choose from the list

Different options display depending on how many compatible fields are in the visual.

5. Toggle **Line** to turn line styling on or off.

You can customize the following line options:

- The weight or thickness of the line.

- The style of the line: solid, dashed, or dotted.
 - The color of the line.
 - The type of line that it is: Linear, Smooth, or Stepped.
6. Toggle **Marker** to turn marker styling on or off.

You can customize the following marker options:

- The weight or thickness of the marker.
 - The style of the marker: circle, triangle, square, diamond, and so on.
 - The color of the marker.
7. For **Axis**, choose whether to display the axis on the left or the right.
 8. Your changes are saved automatically.
 9. (Optional) To undo customizations, choose one or more of the following options:
 - To undo one change, click the undo arrow at top left. Repeat as needed. There is also a redo arrow.
 - To reset the base style for a data series, select **Base style** and then click **Reset to default**.
 - To remove all styling from a data series, listed in **Styled series**, select a field and then click **Remove styling**.

Missing data on visual types in Quick

You can customize how missing data points are visualized in your line charts and area charts. You can choose to have your missing data points appear in the following formats:

- *Broken line*: A disjointed line that breaks when a data point is missing. This is the default missing data format.
- *Continuous line*: Displays a continuous line by skipping over the missing data point and connecting the line to the next available data point in the series. To show a continuous line, the **Show date gaps** box on the **X axis** pane should be unchecked.
- *Show as zero*: Sets the value of the missing data point to zero.


To customize a visual's missing data settings

1. On the analysis page, choose the visual that you want to format.

2. Choose the **Format visual** icon in the upper right corner of the visual to access the **Format visual** menu.
3. Open the **Y axis** pane of the format visual menu and navigate to the **Missing data** section.
4. Select the missing data format that you want.

Reference lines on visuals types in Quick

Reference lines are visual markings in a visual, similar to ruler lines. You typically use a reference line for a value that needs to be displayed with the data. You use the reference line to communicate thresholds or limits in values. The reference line isn't part of the data that's used to build a chart. Instead, it's based on a value that you enter or a field that you identify in the dataset used by a chart.



Quick supports reference lines in the following:

- Bar charts
- Line charts
- Combo charts

You can create, change, and delete reference lines while designing an analysis. You can customize the line pattern, the label font, and the colors for each of those separately. You can show numeric values as numbers, currency, or percent. You can also customize a value's numerical format in the same way that you can customize a field in the field well.

There are two types of reference lines:

- A *constant line* displays at a position that's based on a value that you specify in the format settings. This value doesn't need to relate to any field. You can customize the formatting of the line.
- A *calculated line* displays at a position that's based on a value that is the result of a function. During configuration, you specify which measure (metric) you want to use and which aggregation to apply. These are the same aggregations you can apply to in the field wells. Then, you need to provide an aggregation to apply to the field calculation for the reference line, for example average, minimum, maximum, or percentile. The field needs to be in the dataset used by the chart, although it doesn't need to be displayed in the chart's field wells.

Calculated reference lines aren't supported in 100% stacked charts.

To add or edit a reference line (console)

1. Open the [Quick console](#).
2. Open the analysis that you want to change.
3. Choose the visual that you want to change and open the **Properties** menu.
4. In the **Properties** pane that opens, open the **Reference lines** dropdown, and then choose **ADD NEW LINE**.
5. The **New reference line** menu opens. Use this menu to configure your new reference line. The list below describes all reference line properties that can be configured.
 - **Data**
 - **Type** – The type of reference line that you want to use. Choose one of the following options:
 - To create a constant line based on a single value that you enter, choose **Constant line**.
 - To create a calculated line based on a field, choose **Calculated line**.

- **Value** – (For constant lines only) The value that you want to use. This becomes the location of the line on the visual. It appears immediately, so you can experiment with the setting.
- **Column** – (For calculated lines only) The column that you want to use for the reference line.
- **Aggregated as (column)** – (For calculated lines only) The aggregation that you want to apply to the selected column.
- **Calculate** – (For calculated lines only) The calculation that you want to apply to the aggregation.
- **Percentile value** – (Only if you set **Calculate** to **Percentile**) Enter a number from 1 through 100.
- **Chart type** – (For combo charts) Choose **Bars** or **Lines**.
- **Line style**
 - **Pattern** – The pattern used for the line. Valid options include **Dashed**, **Dotted**, and **Solid**.
 - **Color** – The color used for the line.
- **Label**
 - **Type** – The type of label to display. Valid options include **Value only**, **Custom text**, **Custom text and value**, **No label**. If you choose an option that includes custom text, enter the label text that you want to appear on the line.
 - **Enter custom text (text box)** – (Only if you set **Type** to **Custom text and value**) Choose where to show the value in relation to the label. Valid options are **Left** or **Right**.
 - **Position** – The position of the label in relation to the line. Valid options include a combination of the following: left, middle, right, above, and below.
 - **Value format** – The format to use for the value. Choose one of the following:
 - **Same as value** – Uses the formatting that's already selected for this field in the visualization.
 - **Show as** – Choose from the available options, for example number, currency, or percent.
 - **Format** – Choose from the available formatting options.
 - **Font size** – The font size to use for the label text.
 - **Color** – The color to use for the label text.

6. Choose **Done** to save your selections.

To list existing reference lines

1. Choose the visual that you want to change and open the **Properties** pane.
2. In the **Properties** pane, open the **Reference lines** dropdown, and then choose the ellipsis (three dots) next to the line that you want to change.
3. Choose **Edit**.
4. The **New reference line** menu opens. Use this menu to make changes to your reference line. When you are finished, choose **Done**.

To disable a reference line

1. Choose the visual that you want to change and open the **Properties** pane.
2. In the **Properties** pane, open the **Reference lines** dropdown, and then choose the ellipsis (three dots) next to the line that you want to change.
3. Choose **Disable**.

To delete a reference line

1. Choose the visual that you want to change and open the **Properties** pane.
2. In the **Properties** pane, open the **Reference lines** dropdown, and then choose the ellipsis (three dots) next to the line that you want to change.
3. Choose **Delete**.

Formatting radar charts in Quick

You can customize radar charts in Amazon Quick to arrange your data the way that you want. You can customize the series style, start angle, fill area, and grid shape of a radar chart.

To set the series style of a radar chart

1. Choose the radar chart visual that you want to change, and choose the **Format visual** icon on the top right corner of the visual.
2. In the **Properties** pane on the left, open the **Radar chart** dropdown list.
3. Under **Series style**, choose the style that you want. You can choose between the following styles:

- **LINE.** When selected, the polygons that are created by the data are outlined.
- **AREA.** When selected, the polygons that are created by the data are filled in.

The default selected value is **LINE**.

To choose the start angle of a radar chart

1. Choose the radar chart visual that you want to change, and choose the **Format visual** icon on the top right corner of the visual.
2. In the **Properties** pane on the left, open the **Radar chart** dropdown list.
3. Under **Start angle**, enter the start angle value that you want. The default value is 90 degrees.

To set the fill area of a radar chart

1. Choose the radar chart visual that you want to change, and choose the **Format visual** icon on the top right corner of the visual.
2. In the **Properties** pane on the left, open the **Axis** dropdown list.
3. Select the **Fill grid lines** check box.
4. (Optional) Select colors for the even and odd numbered grid lines.
 - Choose the **Even color** icon that appears, and then choose the color that you want the even numbered grid lines to be. The default color for this value is white.
 - Choose the **Odd color** icon that appears, and then choose the color that you want the odd numbered grid lines to be. The default color for this value is white.

To choose the grid shape of a radar chart

1. Choose the radar chart visual that you want to change, and choose the **Format visual** icon on the top right corner of the visual.
2. In the **Properties** pane on the left, open the **Radar chart** dropdown list.
3. Under **Grid shape**, choose the shape that you want the radar chart grid to be. You can choose between a **POLYGON** and a **CIRCLE**.

Range and scale on visual types in Quick

To change the scale of the values shown on the visual, you can use the **Properties** pane to set the range for one or both axes of the visual. This option is available for the value axes on bar charts, combo charts, line charts, and scatter plots.

By default, the axis range starts at 0 and ends with the highest value for the measure being displayed. For the group-by axis, you can use the data zoom tool on the visual to dynamically adjust the scale.

To set the axis range for a visual

1. On the analysis page, choose the visual that you want to format.
2. Choose the control menu at the upper-right corner of the visual, and then choose the cog icon.
3. On the **Properties** pane, choose **X-Axis** or **Y-Axis**, depending on what type of visual you are customizing. This is the **X-Axis** section for horizontal bar charts, the **Y-Axis** section for vertical bar charts and line charts, and both axes are available for scatter plots. On combo charts, use **Bars** and **Lines** instead.
4. Enter a new name in the box to rename the axis. To revert to the default name, delete your entry.
5. Set the range for the axis by choosing one of the following options:
 - Choose **Auto (starting at 0)** to have the range start at 0 and end around the highest value for the measure being displayed.
 - Choose **Auto (based on data range)** to have the range start at the lowest value for the measure being displayed and end around the highest value for the measure being displayed.
 - Choose **Custom** to have the range start and end at values that you specify.

If you choose **Custom**, enter the start and end values in the fields in that section. Typically, you use integers for the range values. For stacked 100 percent bar charts, use a decimal value to indicate the percentage that you want. For example, if you want the range to be 0–30 percent instead of 0–100 percent, enter 0 for the start value and .3 for the end value.

6. For **Scale**, the default is linear scale. To show logarithmic scale, also called log scale, enable the logarithmic option. Quick chooses the axis labels to display based on the range of values in that axis.

- On a linear scale, the axis labels are evenly spaced to show the arithmetical difference between them. The labels display the numbers in sets like {1000, 2000, 3000...} or {0, 50 million, 100 million...}, but not {10 thousand, 1 million, 1 billion...}.

Use a *linear scale* for the following cases:

- All the numbers that display on the chart are in the same order of magnitude.
- You want the axis labels to be evenly spaced.
- The axis values have a similar number of digits, for example 100, 200, 300, and so on.
- The rate of change between numbers is relatively slow and steady—in other words, your trend line never approaches becoming vertical.

Examples:

- Profits in different regions of the same country
- Costs incurred for manufacture of an item
- On a *logarithmic scale*, the axis values are spaced to show the orders of magnitude as a way of comparing them. The log scale is often used to display very large ranges of values or percentages, or to show exponential growth.

Use logarithmic scale for the following cases:

- The numbers that display on the chart aren't in the same order of magnitude.
- You want the axis labels to be flexibly spaced to reflect the wide range of values in that axis. This might mean that the axis values have a different number of digits, for example 10, 100, 1000, and so on. It might also mean that the axis labels are unevenly spaced.
- The rate of change between numbers is growing exponentially or is too large to display in a meaningful way.
- The customer of your chart understands how to interpret data on a log scale.
- The chart displays values that growing faster and faster. Moving given distance on the scale means the number has been multiplied by another number.

Examples:

- High yield stock prices over a long range of time
- Growth of pandemic infection rates

7. To customize the number of values to show on the axis labels, enter in an integer between 1 and 50.

8. For combo charts, choose **Single Y Axis** to synchronize the Y-axes for both bars and lines into a single axis.
9. Close the **Properties** pane by choosing the **X** icon in the upper-right corner of the pane.

Small multiples axis options

You can configure the x and y axes for each individual panel of a small multiples visual. You can group your data along an independent x-axis or an independent y-axis. You can also position the x and y axes inside or outside the chart to improve the readability of your data.

For small multiples visuals that use an independent x-axis, only the values that are relevant to each panel are shown on the axis. For example, say you have a small multiples visual that uses one panel to represent each region of the United States. With an independent x-axis, each panel only shows states in the region that the panel represents and hides states that are outside of the panel's region.

For small multiples visuals that use an independent y-axis, each panel uses its own y-axis scale that is determined by the range of the data it contains. By default, data labels appear on the inside of the panel.

To configure independent axes for small multiples visuals

1. Select the small multiples visual that you want to change and open the **Format visual** menu.
2. In the **Properties** pane that appears, open the **Multiples options** menu.
3. For **X-axis**, choose **Independent** from the dropdown.

Or, for **Y-axis**, choose **Independent** from the dropdown.

You can revert your changes by choosing **Shared** from the **X-axis** or **Y-axis** dropdown menus.

You can also configure the label positions of the x and y axes of all panels in a small multiples visual. You can choose to display axis labels inside or outside the panel.

To configure the axis label position for small multiples visuals

1. Select the small multiples visual that you want to change and open the **Format visual** menu.
2. In the **Properties** pane that appears, open the **Multiples options** menu.
3. For **X-axis labels**, choose **Inside** or **Outside** from the dropdown.

Or, for **Y-axis labels**, choose **Inside** or **Outside** from the dropdown.

Titles and subtitles on visual types in Quick

In Quick, you can format visual titles and subtitles to meet your business needs. Quick offers rich text formatting for titles and subtitles, and the ability to add hyperlinks and parameters in titles. You can edit titles in the Properties pane, or by double-clicking on a title or subtitle in the visual.

Use the following procedure to customize the way the title and subtitle of a visual is displayed. The visual title is shown by default. After subtitles are created, they're also shown by default.

1. Sign in to Quick at <https://quicksight.aws.amazon.com/>.
2. Open the analysis that you want to update.
3. On the analysis page, choose the visual that you want to format.
4. At the visual's right, choose the **Properties** icon.
5. In the **Properties** pane that opens, choose the **Display settings** tab.
6. To edit the title or subtitle of a visual, choose the paintbrush icon next to **Edit title** or **Edit subtitle**. Alternatively, you can choose the eyeball icon next to **Edit title** or **Edit subtitle** to hide the title or subtitle, shown in the following image.
7. In the **Edit title** or **Edit subtitle** popup that opens, you can use the following options to make the updates that you want:
 - To enter a custom title or subtitle, enter your title or subtitle text in the editor. Titles can be up to 120 characters long, including spaces. Subtitles can be up to 500 characters long.
 - To change the font type, choose a font type from the list at left.
 - To change the font size, choose a size from the list at right.
 - To change the font weight and emphasis, or to underline or strikethrough text, choose the bold, emphasis, underline, or strikethrough icons.
 - To change the font color, choose the color (Abc) icon, and then pick a color. You can also enter a hexadecimal number or RGB values.
 - To add an unordered list, choose the unordered list icon.
 - To change the text alignment, choose the left, center, or right alignment icons.

- To add a parameter to a title or subtitle, choose an existing parameter from the list under **Parameters** at right. For more information about how to create parameters, see [Setting up parameters in Amazon Quick](#).
- To add a hyperlink, highlight the text that you want to link, choose the hyperlink icon, and then choose from the following options:
 - For **Enter link**, enter the URL that you want to link to.

Choose the + icon at right to add an existing parameter, function, or computation to the URL.

- To edit the display text, enter text for **Display text**.
- To open the hyperlink in the same browser tab as Quick, select **Same tab**.
- To open the hyperlink in a new browser tab, select **New tab**.
- To delete the hyperlink, choose the delete icon at bottom left.

When finished configuring the hyperlink, choose **Save**.

8. When you are finished, choose **Save**.
9. For **Alt text**, enter the alt text that you want for the visual.
10. When you are finished, close the properties pane.

Tooltips on visual types in Quick

When you hover your cursor over any graphical element in an Quick visual, a tooltip appears with information about that specific element. For example, when you hover your cursor over dates in a line chart, a tooltip appears with information about those dates. By default, the fields in the Fields well determine what information displays in tooltips. Tooltips can display up to 10 fields.

You can provide your viewers with additional information about data in your visual, customizing what viewers can see. You can even prevent tooltips from appearing when viewers hover a cursor over an element. To do this, you can customize the tooltips for that visual.

Customizing tooltips in a visual

Use the following procedure to customize tooltips in a visual.

To customize tooltips in a visual

1. On the analysis page, choose the visual that you want to format.

2. On the menu in the upper-right corner of the visual, choose the **Format visual** icon.
3. In the **Properties** pane that opens, choose **Tooltip**.
4. For **Type**, choose **Detailed tooltip**. A new set of options appear.

To show or hide titles in a tooltip

- Choose **Use primary value as title**.

Clearing the option hides titles in the tooltip. Selecting the option shows the primary field value as the title in the tooltip.

To show or hide aggregations for fields in the tooltip

- Choose **Show aggregations**.

Clearing the option hides the aggregation for fields in the tooltip. Selecting the option shows the aggregation for fields in the tooltip.

To add a field to the tooltip

1. Choose **Add field**.
2. In the **Add field to tooltip** page that opens, choose **Select field** and then select a field from the list.

You can add up to 10 fields to tooltips.

3. (Optional) For **Label**, enter a label for the field. This option creates a custom label for the field in the tooltip.
4. (Optional) Depending on whether you add a dimension or a measure, choose how you want the aggregation to display in the tooltip. If you don't select an option, Quick uses the default aggregation.

If you add a measure to the tooltip, you can select how you want the field to be aggregated. To do so, choose **Select aggregation**, and then select an aggregation from the list. For more information about the types of aggregations in Quick, see [Changing field aggregation](#).

5. Choose **Save**.

A new field is added to the list of fields in your tooltip.

To remove a field from the tooltip

- Under the **Fields** list, select the field menu for the field that you want to remove (the three dots) and choose **Hide**.

To rearrange the order of the fields in the tooltip

- Under the **Fields** list, select the field menu for a field (the three dots) and choose either **Move up** or **Move down**.

To customize the label for a field in the tooltip

- Select the field menu for the field that you want to customize (the three dots) and choose **Edit**.
- In the **Edit tooltip field** page that opens, for **Label**, enter the label that you want to appear in the tooltip.
- Choose **Save**.

Using sheet tooltips in Quick


Sheet tooltips transform how viewers explore data by providing rich context without disrupting their analysis flow. Instead of navigating away from a visual or opening separate sheets, viewers get instant access to detailed breakdowns, trends, and supporting information, making dashboards more intuitive and reducing the need for multiple sheets.

Sheet tooltips are available on interactive sheets only. They are not supported on paginated reports. You can duplicate a tooltip sheet to another tooltip sheet, or duplicate a tooltip sheet to a regular interactive sheet. Additionally, you can duplicate a visual to a tooltip sheet.

How sheet tooltips work

When an author creates a sheet tooltip, a tooltip sheet is created and associated with a visual. This tooltip sheet works like a regular sheet. You can add visuals, text boxes, and images to it using a free-form layout. When a viewer hovers over a data point, the tooltip sheet inherits all filters from the source visual and adds an additional filter for the specific data point. For example, if your source visual is filtered to "2025 data" and a viewer hovers over "Electronics," the tooltip shows Electronics data for 2025 only.

Consider a bar chart showing sales by product category. You could create a sheet tooltip that shows a trend line of monthly sales, a KPI of year-over-year growth, and a text box with the category name, all filtered to whichever category the viewer hovers over.



Sheet tooltip limits

The following limits apply to sheet tooltips:

- Up to 50 tooltip sheets per analysis
- Up to 5 visuals per tooltip sheet
- Up to 5 text boxes per tooltip sheet
- Up to 5 images per tooltip sheet
- Tooltip sheets use free-form layout only
- Layer map visuals are not allowed on tooltip sheets
- Maximum size of a tooltip sheet is 640px wide by 720px tall


Creating a sheet tooltip

Use the following procedure to create a sheet tooltip for a visual.

To create a tooltip sheet

1. On the analysis page, choose the visual that you want to add a sheet tooltip to.

2. On the menu in the upper-right corner of the visual, choose the **Format visual** icon.
3. In the **Properties** pane that opens, choose **Interactions > Tooltip**.
4. For **Type**, choose **Sheet tooltip**.



5. Choose **Create tooltip sheet**. You will automatically navigate to a tooltip sheet editing experience. A tooltip name is auto-generated and you can edit it by choosing the tab title.
6. Add visuals, text boxes, or images to the tooltip sheet. Arrange them using the free-form layout.
7. When you are finished, return to the source sheet by choosing the **Back** button located to the left of the sheet tooltip title. To preview the tooltip, hover over any data points in the visual.

Assigning a tooltip sheet to a visual

When you select **Sheet tooltip** as the tooltip type in the **Properties** pane, a control appears that lets you select all tooltip sheets available in the analysis. You can assign one tooltip sheet to multiple visuals or create separate tooltip sheets for each visual.


If you would like to apply the same tooltip sheet to another visual, you can do this by assigning one tooltip sheet to multiple visuals in the **Interactions > Tooltip** accordion in the **Properties** pane.

Editing a tooltip sheet

Use the following procedure to edit an existing sheet tooltip.

To edit a tooltip sheet

1. Choose any visual where a sheet tooltip is enabled.
2. Open the **Properties** pane and navigate to **Interactions > Tooltip**.
3. In the **Tooltip** accordion, select the tooltip that you would like to edit and choose the edit icon next to the tooltip sheet name to navigate to it.
4. Make your changes to the visuals, text boxes, or images on the tooltip sheet.



Switching between tooltip types

You can switch a visual's tooltip between basic, detailed, and sheet tooltip types at any time.

To change the tooltip type

1. Choose the visual that you want to update.
2. Open the **Properties** pane, choose **Interactions**, and then choose **Tooltip**.
3. For **Type**, select the tooltip type that you want: **Basic tooltip**, **Detailed tooltip**, or **Sheet tooltip**.

Note

Switching away from a sheet tooltip preserves your work. You can always switch back without losing your tooltip sheet design.

Sheet tooltip considerations

Keep the following in mind when working with sheet tooltips:

- Tables and pivot tables support sheet tooltips but not basic or detailed tooltips.
- Visuals in a tooltip sheet do not support context menus, on-visual menus, or custom actions.
- [Using custom actions for filtering and navigating](#) on visuals in a tooltip sheet are not supported when the sheet is rendered as a tooltip.
- Sheet tooltips support filters, cross-sheet filtering, and parameters. Filter controls are not supported.
- Sheet descriptions are not displayed on tooltip sheets.
- Cross-sheet filters cannot be scoped to tooltip sheets.
- An analysis must contain at least one regular interactive sheet. An analysis cannot consist of only tooltip sheets.
- Layer map visuals cannot be placed inside a tooltip sheet.
- Tooltips on tooltip sheets are not supported.
- Sheet tooltips are not supported on the following chart types: Sankey, Waterfall, KPI, Radar, Wordcloud, Custom content, and Highcharts.

These limits ensure tooltip sheets load quickly and maintain a focused, scannable experience for viewers. For more complex analysis, consider using drill-down actions or separate detail sheets.

Hiding tooltips in a visual

If you don't want tooltips to appear when you hover your cursor over data in a visual, you can hide them.

To hide tooltips in a visual

1. On the analysis page, choose the visual that you want to format.

2. On the menu in the upper-right corner of the visual, choose the **Format visual** icon.
3. In the **Properties** pane that opens, choose **Tooltip**.
4. Choose **Show tooltip**.

Clearing the option hides tooltips for the visual. Selecting the option shows them.

Customizing data presentation

To gain further insight into your data when creating visuals (charts) in an Quick analysis, you can sort and filter data in a visual. You can also change the granularity of date fields, data type, role, and format of fields in a visual.

Topics

- [Changing fields used by a visual in Amazon Quick](#)
- [Sorting visual data in Amazon Quick](#)

Changing fields used by a visual in Amazon Quick

You can add or modify fields for a visual by using the **Fields list** pane, the field wells, or the on-visual editors or drop targets on the visual.

The field wells, on-visual editors, and drop targets available for a specific visual depends on the visual type selected. For details, see the appropriate visual type topic in the [Visual types in Amazon Quick Sight](#) section.

Important

You can also change the data type and format of numeric fields by using field wells and on-visual editors. If you change a field in this way, it changes for the selected visual only. For more information about changing numeric field data types and formats, see [Changing fields used by a visual in Amazon Quick](#).

Use the following topics to learn more about adding, removing, and modifying fields on a visual.

Topics

- [Using visual field controls](#)

- [Adding or removing a field](#)
- [Changing the field associated with a visual element](#)
- [Changing field aggregation](#)
- [Changing date field granularity](#)
- [Customizing a field format](#)

Using visual field controls

You can edit the fields used by a visual with user interface (UI) controls.

You can use these controls as follows:

- Create a visual and assign fields to different elements on it by selecting fields in the **Fields list** pane, or dragging fields to field wells or drop targets.
- Change the field associated with a visual element by dragging a field to a drop target or field well, or selecting a different field in a field well or on-visual editor.
- Change field aggregation or date granularity by using the field wells or the on-visual editors.

The field wells, on-visual editors, and drop targets available on a specific visual depends on the visual type selected.

Dragging fields to drop targets or field wells

When you drag a field to either a drop target or field well, Amazon Quick provides you with information about whether the target element expects a measure or a dimension. Amazon Quick also provides you with information about whether that element is available for field assignment.

For example, when you drag a measure to the value drop target on a new single-measure line chart, you see the drop target color-coded green. That green color coding indicates that the drop target expects a measure. The drag label indicates that the target is available to add a field.

When you drag a dimension to the x-axis or color drop target on a new line chart, you see a label color-coded blue. That blue color coding indicates that the drop target expects a dimension. The drag label indicates that the target is available to add a field.

You can also drag a measure or dimension to a drop target on a line chart where the element is already associated with a field. In this case, the drag label indicates that you are replacing the field currently associated with the drop target.

Adding or removing a field

You can add a field to a visual by choosing it on the **Fields list** pane. You can also drag it to a drop target on the visual or to a field well. There is a 1:1 correspondence of drop targets to field wells for each visual type, so you can use either method.

On some charts, the **Axis title** field is hidden when there are two or more fields in the **Value** field on any side of the chart. This effect can happen with the following charts:

- Bar charts
- Line charts
- Box plots
- Combo charts
- Waterfall charts

To remove a field from a visual, clear selection from it in the **Fields list** pane. Or choose an on-visual editor or field well that uses that field, and then choose **Remove** from the context (right-click) menu.

Adding a field by selecting it in the fields list pane

You can also let Amazon Quick map the field to the most appropriate visual element. To do so, choose the field in the **Fields list** pane. Amazon Quick adds the field to the visual by populating the first empty field well that corresponds with that field type (either measure or dimension). If all of the visual elements are already populated, Amazon Quick determines the most appropriate field well and replaces the field in it with the field you selected.

Adding a field by using a drop target

To add a field to a visual by using a drop target, first choose a field in the **Fields list** pane. Then drag the field to your chosen drop target on the visual, making sure the drop indicator shows that the field is being added.

Adding a field by using a field well

To add a field to a visual by using a field well, choose a field in the **Fields list** pane. Then drag the field to the target field well, making sure that the drop indicator shows that the field is being added.

1. Drag a field item into a **Field well**.

2. Drag the field that you want to add from the **Fields list** pane to the appropriate field well.

Note

You can add the same value to the same visual multiple times. You can do so to show the same value with different aggregations or table calculations applied. By default, the fields all display the same label. You can edit the names by using the **Properties** panel, which you open by choosing the **V**-shaped icon at top right.

Changing the field associated with a visual element

You can change the field assigned to an element in a visual by using the field wells, drop targets, or the on-visual editors on the visual. For pivot tables, use field wells or drop targets because this visual type doesn't provide on-visual editors.

Change a field mapping by using an on-visual editor

Use the following procedure to modify the mapping of a field to a visual element.

To modify the mapping of a field by using an on-visual editor

1. On the visual, choose the on-visual editor for the visual element for which you want to change the field.
2. On the on-visual editor menu, choose the field that you want to associate with that visual element.

Changing a field mapping by using a drop target

To modify the mapping of a field to a visual element by using a drop target, choose a field in the **Fields list** pane. Then drag the field to a drop target on the visual, making sure that the drop indicator shows that the field is being replaced.

Changing a field mapping by using a field well

Use the following procedure to modify the mapping of a field to a visual element.

To modify the mapping of a field by using a field well

1. Drag a field item into a **Field well**.

2. Choose the field well that represents the element that you want to remap, and then choose a new field from the menu that appears.

Changing field aggregation

You can apply functions to fields to display aggregate information, like the sum of the sales for a given product. You can apply an aggregate function by using the options in either an on-visual editor or a field well. The following aggregate functions are available in Amazon Quick:

- Average – Calculates the average value for the selected field.
- Count – Provides a count of the number of records containing the selected measure for a given dimension. An example is a count of Order ID by State.
- Distinct Count – Provides a count of how many different values are in the selected measure, for the selected dimension or dimensions. An example is a count of Product by Region. A simple count can show how many products are sold for each region. A distinct count can show how many different products are sold for each region. You might have sold 2,000 items, but only two different types of items.
- Max – Calculates the maximum value for the selected field.
- Min – Calculates the minimum value for the selected field.
- Median – Calculates the median value of the specified measure, grouped by the chosen dimension or dimensions.
- Sum – Totals all of the values for the selected field.
- Standard Deviation – Calculates the standard deviation of the set of numbers in the specified measure, grouped by the chosen dimension or dimensions, based on a sample or on a biased population.
- Variance – Calculates the variance of the set of numbers in the specified measure, grouped by the chosen dimension or dimensions, based on a sample or on a biased population.
- Percentile – Computes the *n*th percentile of the specified measure, grouped by the chosen dimension or dimensions.

All aggregate functions can be applied to numeric fields. *Count* is automatically applied to a dimension if you choose to use it in a field well that expects a measure. If you have used a dimension in that way, you can also change the aggregate function applied to it. You can't apply aggregate functions to fields in dimension field wells.

The visual elements that support aggregated fields varies by visual type.

Changing or adding aggregation on a field by using an on-visual editor

Use the following procedure to change or add aggregation on a field.

To change or add aggregation on a field

1. On the visual, choose the on-visual editor for the field that you want to apply aggregation to.
2. On the on-visual editor menu, choose **Aggregate**, then choose the aggregate function that you want to apply.

Changing or adding aggregation to a field by using a field well

Use the following procedure to add aggregation to a field for a pivot table visual.

To add aggregation to a field for a pivot table visual

1. Drag a field item into a **Field well**.
2. Choose the field well containing the field that you want to apply an aggregate function to.
3. On the field well menu, choose **Aggregate**, then choose the aggregate function that you want to apply.

Changing date field granularity

You can change the granularity for a date field on a visual to determine the intervals for which item values are shown. You can set the date field granularity to one of the following values:

- Year
- Quarter
- Month
- Week
- Day (this is the default)
- Hour
- Minute
- Second

Hour and minute are available only if the field contains time data.

Changing date field granularity by using an on-visual editor

Use the following procedure to change date field granularity by using an on-visual editor.

To change date field granularity with an on-visual editor

1. On the visual, choose the field well for the date field whose granularity you want to change.
2. On the field well menu, choose **Aggregate**, then choose the time interval that you want to apply, as shown following:

Changing date field granularity by using a field well

Use the following procedure to change date field granularity by using a field well.

To change date field granularity with a field well

1. Drag a field item into a **Field well**.
2. Choose the field well containing the date field, and then choose **Aggregate**. Choose the date granularity that you want to use.

Customizing a field format

Use the following procedure to customize the appearance of fields in an analysis.

To customize the appearance of fields in an analysis

1. In an analysis, choose a field to format, either by choosing it in the field well or in the **Fields list** of the **Visualize** pane.
2. Choose **Show as** to change how the field shows in the analysis, and choose from the options on the context menu. The list of available options varies based on the field's data type. If you choose a non-numeric field from the fields list, you can change the *count format*, which is the formatting used when the field is counted.
3. Choose **Format** to change the format of the field, and choose from the options on the context menu. If you don't see an option that you want to use, choose **More formatting options** from the context menu.

The **Format Data** pane opens, presenting options for the type of numeric or date field you chose.

The options for **Show as** from the context menu now appear in the drop-down list at the top of the **Format Data** pane. The rest of the options are specific to the data type and how you choose to show the field.

For date and time data, the default format pattern is YYYY-MM-DDTHH:mm:ssZZ, for example 2016-09-22T17:00:00-07:00.

For numbers, you can choose from the following units to display after the number:

- No unit suffix. This is the default.
- Thousands (K)
- Millions (M)
- Billions (B)
- Trillions (T)
- A custom unit prefix or suffix

For currency, you can choose from the following symbols:

- Dollars (\$)
- Euros (€)
- Pounds (£)
- Yen (¥)

Changing a field format

You can change the format of a field within the context of an analysis. The formatting options available for fields vary based on the field's data type.

Use menu options in the **Field list** pane or the visual field wells to make simple format changes, or use the **Format data** pane to make more extensive formatting changes.

Topics

- [Format a currency field](#)
- [Format a date field](#)
- [Format a number field](#)
- [Format a percent field](#)
- [Format a text field](#)
- [Return a field's format to default settings](#)

Format a currency field

When you format a currency field, you can either choose the currency symbol from a list of common options, or open the **Format data** pane and manually format the field. Manually formatting the field allows you to choose which symbol to use, which separators to use, the number of decimal places to show, which units to use, and how to display negative numbers.

Changing a field format changes it for all visuals in the analysis, but does not change it in the underlying dataset.

If you want to choose the symbol for a currency field from a list of common options, you can access such a list in several ways. You can access it from the **Field list** pane, an on-visual editor, or a visual field well.

To select a currency field's symbol by choosing a list option

1. Choose one of the following options:
 - In the **Field list** pane, choose the selector icon to the right of the number field that you want to format.
 - On any visual that contains an on-visual editor associated with the currency field that you want to format, choose that on-visual editor. Expand the **Field wells** pane, and then choose the field well associated with the number field that you want to change.
2. Choose **Format**, and then choose the currency field that you want:
 - Display in dollars (\$).
 - Display in pounds (£).
 - Display in euros (€).
 - Display in yen or yuan (¥).

To manually change a currency field's format

1. Choose one of the following options:
 - In the **Field list** pane, choose the selector icon to the right of the number field that you want to format.
 - On any visual that contains an on-visual editor associated with the number field that you want to format, choose that on-visual editor. Expand the **Field wells** pane, and then choose the field well associated with the number field that you want to change.
2. Choose **Format**, and then choose **More Formatting Options**.

The **Format data** pane opens.

3. Expand the **Symbol** section and choose from the following options:
 - Display in dollars (\$). This is the default.
 - Display in pounds (£).
 - Display in euros (€).
 - Display in yen or yuan (¥).
4. Expand the **Separators** section and choose from the following options:
 - Under **Decimal**, choose a dot or a comma for the decimal separator. A dot is the default. If you choose a comma instead, use a dot or a space as the thousands separator.
 - Under **Thousands**, select or clear **Enabled** to indicate whether you want to use a thousands separator. **Enabled** is selected by default.
 - If you are using a thousands separator, choose whether to use a comma, dot, or space for the separator. A comma is the default. If you choose a dot instead, use a comma as the decimal separator.
5. Expand the **Decimal Places** section and choose the number of decimal places to use. The default is 2. Field values are rounded to the decimal places specified. For example, if you specify two decimal places, the value 6.728 is rounded to 6.73.
6. Expand the **Units** section and choose from the following options:
 - Choose the unit to use. Choosing a unit adds the appropriate suffix to the number value. For example, if you choose **Thousands**, a field value of 1234 displays as 1.234K.

The unit options are as follows:

- No unit suffix. This is the default.
 - Thousands (K)
 - Millions (M)
 - Billions (B)
 - Trillions (T)
 - If you want to use a custom prefix or suffix, specify it in the **Prefix** or **Suffix** box. Using a custom suffix is a good way to specify a currency suffix outside of those already offered by Amazon Quick. You can specify both. You can also specify a custom prefix in addition to the suffix added by selecting a unit.
7. Expand the **Negatives** section and choose whether to display a negative value by using a minus sign or by enclosing it in parentheses. Using a minus sign is the default.
 8. Expand the **Null values** section and choose whether to display null values as `null` or as a custom value. Using `null` is the default.

 Note

When using a table or pivot table, null values only display for fields that are placed in the **Rows**, **Columns**, or **Group by** field wells. Null values for fields in the **Values** field well appear empty in the table or pivot table.

Format a date field

When you format a date field, you can choose a list of common formatting options. Or you can open the **Format data** pane to choose from a list of common formats, or specify custom formatting for the date and time values.

Changing a field format changes it for all visuals in the analysis that use that dataset, but does not change it in the dataset itself.

If you want to format a date field by choosing from a list of common options, you can access such a list in several ways. You can access it from the **Field list** pane, a visual on-visual editor, or a visual field well.

To change a date field's format by choosing a list option

1. Choose one of the following options:

- In the **Field list** pane, choose the selector icon to the right of the number field that you want to format.
 - On any visual that contains an on-visual editor associated with the number field that you want to format, choose that on-visual editor. Expand the **Field wells** pane, and then choose the field well associated with the number field that you want to change.
2. Choose **Format**, and then choose the format that you want. The following quick formatting options are offered for date fields:
 - Show the month, day, year, and time.
 - Show the month, day, and year.
 - Show the month and year.
 - Show the year.

To manually change a date field's format

1. Choose one of the following options:
 - In the **Field list** pane, choose the selector icon to the right of the number field that you want to format.
 - On any visual that contains an on-visual editor associated with the number field that you want to format, choose that on-visual editor. Expand the **Field wells** pane, and then choose the field well associated with the number field that you want to change.

2. Choose **Format**, and then choose **More Formatting Options**.

The **Format data** pane opens.

3. Expand the **Date** section. Choose an existing date format, or choose **Custom** and specify a format pattern in the **Custom** section lower down in the **Format data** pane. If you choose **Custom** for the **Date** section, you must also choose **Custom** for the following **Time** section. The pattern that you specify in the **Custom** section must include any date and time formatting that you want.

The default selection is **Custom**, with a default format pattern of MMM D, YYYY h:mm, for example Sep 20, 2022 5:30pm.

4. Expand the **Time** section. Choose an existing time format, or choose **Custom** and specify a format pattern in the **Custom** section lower down in the **Format data** pane. If you choose

Custom for the **Time** section, you must also choose **Custom** for the preceding **Date** section. The pattern that you specify in the **Custom** section must include any date and time formatting that you want.

The default selection is **Custom**, with a default format pattern of MMM D, YYYY h:mma, for example Sep 20, 2022 5:30pm.

5. If you chose **Custom** in the **Date** and **Time** sections, expand the **Custom** section and specify the format pattern that you want, using the format pattern syntax specified in [Moment.js Display Format](#) in the Moment.js JavaScript documentation.

 Note

The time zone related display token Z from the Moment.js library is supported in Quick, but the z token is not.

If you chose something other than **Custom** in the **Date** and **Time** sections, **Custom** is populated with the format pattern that reflects your selections. For example, if you chose Jun 21, 2016 in the **Date** section and 17:00:00pm in the **Time** section, the **Custom** section shows the format pattern MMM D, YYYY H:mm:ssa.

6. (Optional) Expand the **Custom** section and use **Preview** to verify your specified format.
7. Expand the **Null values** section and choose whether to display null values as `null` or as a custom value. Using `null` is the default.

Customizing date formats in Quick

In Quick, you can customize how dates are formatted in your filter and parameter controls. For example, you can specify to format the date in a control as 20-09-2021, or, if you'd rather, as 09-20-2021. You can also specify to shorten the month in your dates (such as September) to three letters (Sep), among other customizations.

Following is a list of tokens you can use to create custom date formats. You can use them in combination with one another to control how dates appear in your controls.

List of supported tokens for formatting dates

Use the following tokens to customize the format of dates in Quick.

Example	Description	Token
0–6	Numeric representation of a particular day of the week. 0 is Sunday and 6 is Saturday.	d
Mo–Su	A 2-character textual representation of a particular day of the week.	dd
Mon–Sun	A 3-character textual representation of a particular day of the week.	ddd
Monday–Sunday	A textual representation of a particular day of the week.	dddd
99 or 21	A 2-digit representation of a year.	YY
1999 or 2021	A full, 4-digit numeric representation of a year.	YYYY
1–12	Number of a month, without leading zeros.	M
1st, 2nd, to 12th	Number of a month without leading zeros and with an ordinal suffix.	Mo
01–12	Number of a month with leading zeros.	MM

Example	Description	Token
Jan–Dec	A 3-digit textual representation of a month.	MMM
January–December	A full textual representation of a month.	MMMM
1–4	A numeric representation of a quarter.	Q
1st–4th	A numeric representation of a quarter with an ordinal suffix.	Qo
1–31	Day of the month without leading zeros.	D
1st, 2nd, to 31st	Day of the month without leading zeros and an ordinal suffix.	Do
01–31	A 2-digit day of the month with leading zeros.	DD
1–365	Day of the year without leading zeros.	DDD
001–365	Day of the year with leading zeros.	DDDD
1–53	Week of the year without leading zeros.	w

Example	Description	Token
1st–53rd	The week of the year without leading zeros and with an ordinal suffix.	wo
01–53rd	Week of the year with leading zeros.	ww
1–23	Hours, in 24-hour format, without leading zeros.	H
01–23	Hours, in 24-hour format, with leading zeros.	HH
1–12	Hours, in 12-hour format, without leading zeros.	h
01–12	Hours, in 12-hour format, with leading zeros.	hh
0–59	Minutes without leading zeros.	m
00–59	Minutes with leading zeros.	mm
0–59	Seconds without leading zeros.	s
00–59	Seconds with leading zeros.	ss
am or pm	am/pm	a

Example	Description	Token
AM or PM	AM/PM	A
1632184215	Unix timestamp.	X
1632184215000	Millisecond Unix timestamp.	x
Z	Zero UTC offset.	Z

The following date types are not supported.

- Time zones offset with a colon. For example, +07:00.
- Time zones offset without a colon. For example, +0730.

Preset date formats

To quickly customize dates and times to appear as one of the following example formats, you can use the following Quick preset tokens.

Example	Token
8:30 PM	LT
8:30:25 PM	LTS
August 2 1985	LL
Aug 2 1985	11
August 2 1985 08:30 PM	LLL

Example	Token
Aug 2 1985 08:30 PM	l11
Thursday, August 2 1985 08:30 PM	LLLL
Thu, Aug 2 1985 08:30 PM	l111

Common date formats

Following are three common date examples and their associated token formats for your quick reference.

Example	Token Format
Sep 20, 2021	MMM DD, YYYY
20-09-21 5pm	DD-MM-YY ha
Monday, September 20, 2021 17:30:15	dddd, MMMM DD, YYYY HH:mm:ss

Adding words to dates

To include words in your date formats, such as the word "of" in *20th of Sep, 2021*, enter backslashes (\) before each character in the word. For example, for the 20th of Sep, 2021 date example, use the following token format: Do \o\f MMM, YYYY.

Example: Customizing the date format in a filter control

Use the following procedure to learn how to use date token formats to customize dates for a filter control.

To learn to customize dates for a filter control with data tokens

1. In a Quick analysis, choose the filter control that you want to customize.
2. On the filter control, choose the **Edit control** icon.

3. On the **Edit control** page that opens, for **Date format**, enter the custom date format that you want. Use the tokens listed previously in this topic.

For example, let's say that you want to customize your dates using the following format: *Sep 3rd, 2020 at 5pm*. To do so, you can enter the following token format:

MMM Do, YYYY \a\t ha

A preview of the date format appears below the input field as you enter each token.

4. Choose **Apply**.

The dates in the control update to the format you specified.

Format a number field

When you format a number field, you can choose the decimal place and thousand separator format from a list of common options. Or you can open the **Format Data** pane and manually format the field. Manually formatting the field enables you to choose which separators to use and the number of decimal places to show. It also enables you to choose which units to use, and how to display negative numbers.

Changing a field format changes it for all visuals in the analysis, but does not change it in the underlying dataset.

If you want to format a number field by choosing from a list of common options, you can access such a list from the **Field list** pane, an on-visual editor, or a visual field well.

To change a number field's format by choosing a list option:

- Choose one of the following options:
 - In the **Field list** pane, choose the selector icon to the right of the number field that you want to format.
 - On any visual that contains an on-visual editor associated with the number field that you want to format, choose that on-visual editor. Expand the **Field wells** pane, and then choose the field well associated with the number field that you want to change.
- Choose **Format**, and then choose the format that you want. The following quick formatting options are offered for number fields:
 - Use commas to separate groups of thousands and use a decimal point to show the fractional part of the number, for example 1,234.56.

- Use a decimal point to show the fractional part of the number, for example 1234.56.
- Show the number as an integer and use commas to separate groups of thousands, for example 1,234.
- Show the number as an integer, for example 1234.

To manually change a number field's format:

1. Choose one of the following options:

- In the **Field list** pane, choose the selector icon to the right of the number field that you want to format.
- On any visual that contains an on-visual editor associated with the number field that you want to format, choose that on-visual editor. Expand the **Field wells** pane, and then choose the field well associated with the number field that you want to change.

2. Choose **Format**, and then choose **More Formatting Options**.

The **Format data** pane opens.

3. Expand the **Separators** section and choose from the following options:

- Under **Decimal**, choose a dot or a comma for the decimal separator. A dot is the default. If you choose a comma instead, use a dot or a space as the thousands separator.
- Under **Thousands**, select or clear **Enabled** to indicate whether you want to use a thousands separator. **Enabled** is selected by default.
- If you are using a thousands separator, choose whether to use a comma, dot, or space for the separator. A comma is the default. If you choose a dot instead, use a comma as the decimal separator.

4. Expand the **Decimal Places** section and choose from the following options:

- Choose **Auto** to have Amazon Quick automatically determine the appropriate number of decimal places, or choose **Custom** to specify a number of decimal places. **Auto** is the default.
- If you chose **Custom**, enter the number of decimal places to use. Field values are rounded to the decimal places specified. For example, if you specify two decimal places, the value 6.728 is rounded to 6.73.

5. Expand the **Units** section and choose from the following options:

- Choose the unit to use. Choosing a unit adds the appropriate suffix to the number value. For example, if you choose **Thousands**, a field value of 1234 displays as 1.234K.

The unit options are as follows:

- No unit suffix. This is the default.
 - Thousands (K)
 - Millions (M)
 - Billions (B)
 - Trillions (T)
- If you want to use a custom prefix or suffix, specify it in the **Prefix** or **Suffix** box. You can specify both. You can also specify a custom prefix in addition to the suffix added by selecting a unit.
6. Expand the **Negatives** section and choose whether to display a negative value by using a minus sign or by enclosing it in parentheses. Using a minus sign is the default.
7. Expand the **Null values** section and choose whether to display null values as `null` or as a custom value. Using `null` is the default.

 Note

When using a table or pivot table, null values only display for fields that are placed in the **Rows**, **Columns**, or **Group by** field wells. Null values for fields in the **Values** field well appear empty in the table or pivot table.

Format a percent field

When you format a percent field, you can choose the number of decimal places from a list of common options. Or you can open the **Format data** pane and manually format the field. Manually formatting the field enables you to choose which separators to use. It also enables you to choose the number of decimal places to show and how to display negative numbers.

Changing a field format changes it for all visuals in the analysis, but does not change it in the underlying dataset.

If you want to choose the number of decimal places for a percent field from a list of common options, you can access such a list in several ways. You can access it from the **Field list** pane, an on-visual editor, or a visual field well.

To change a percent field's number of decimal places by choosing a list option

1. Choose one of the following options:
 - In the **Field list** pane, choose the selector icon to the right of the number field that you want to format.
 - On any visual that contains an on-visual editor associated with the percent field that you want to format, choose that on-visual editor. Expand the **Field wells** pane, and then choose the field well associated with the number field that you want to change.
2. Choose **Format**, and then choose the number of decimal places that you want. The following quick formats are offered for percent fields:
 - Display the value with two decimal places.
 - Display the value with one decimal place.
 - Display the value with no decimal places.

To manually change a percent field's format

1. Choose one of the following options:
 - In the **Field list** pane, choose the selector icon to the right of the number field that you want to format.
 - On any visual that contains an on-visual editor associated with the number field that you want to format, choose that on-visual editor. Expand the **Field wells** pane, and then choose the field well associated with the number field that you want to change.

2. Choose **Format**, and then choose **More Formatting Options**.

The **Format data** pane opens.

3. Expand the **Separators** section and choose from the following options:
 - Under **Decimal**, choose a dot or a comma for the decimal separator. A dot is the default. If you choose a comma instead, use a dot or a space as the thousands separator.
 - Under **Thousands**, select or clear **Enabled** to indicate whether you want to use a thousands separator. **Enabled** is selected by default.
 - If you are using a thousands separator, choose whether to use a comma, dot, or space for the separator. A comma is the default. If you choose a dot instead, use a comma as the decimal separator.

4. Expand the **Decimal Places** section and choose from the following options:
 - Choose **Auto** to have Amazon Quick automatically determine the appropriate number of decimal places, or choose **Custom** to specify a number of decimal places. **Auto** is the default.
 - If you chose **Custom**, enter the number of decimal places to use. Field values are rounded to the decimal places specified. For example, if you specify two decimal places, the value 6.728 is rounded to 6.73.
5. Expand the **Negatives** section and choose whether to display a negative value by using a minus sign or by enclosing it in parentheses. Using a minus sign is the default.
6. Expand the **Null values** section and choose whether to display null values as null or as a custom value. Using null is the default.

 Note

When using a table or pivot table, null values only display for fields that are placed in the **Rows**, **Columns**, or **Group by** field wells. Null values for fields in the **Values** field well appear empty in the table or pivot table.

Format a text field

When you format a text field, you can choose how to display null values using the **Field list** pane, an on-visual editor, or a visual field well.

To choose how to display a text field's null values

1. Choose one of the following options:
 - In the **Field list** pane, choose the selector icon to the right of the number field that you want to format.
 - On any visual that contains an on-visual editor associated with the percent field that you want to format, choose that on-visual editor. Expand the **Field wells** pane, and then choose the field well associated with the number field that you want to change.
2. Choose **Format**, and then choose **More Formatting Options**.

The **Format data** pane opens.

3. Expand the **Null values** section and choose whether to display null values as null or as a custom value. Using null is the default.

Return a field's format to default settings

Use the following procedure to return a field's format to the default settings.

To return a field's format to the default settings

1. In the **Field list** pane, choose the selector icon to the right of the field that you want to reset.
2. Choose **Format**, and then choose **More Formatting options**.

The **Format data** pane opens.

3. Choose **Reset to defaults**.

Sorting visual data in Amazon Quick

You can sort data using multiple methods for most visual types. You can choose the sort order of on-visual data by using the quick sort option or field wells. You can also use field wells to sort data by an off-visual metric. The visual element you can sort by depends on the visual type and whether sorting is supported for that visual. For more information on which visual types support sorting, see [Analytics formatting per type in Quick](#).

Pivot tables behave differently than tables when sorting values. For more information about sorting pivot tables, see [Sorting pivot tables in Quick](#).

For SPICE datasets, you can sort text strings of sizes up to the following limitations:

- Up to two million (2,000,000) unique values
- Up to 16 columns

When you exceed the limitations, the visual displays a notification at the upper right.

You can sort any visual type that supports sorting. If a visual type supports sorting, you can sort by using either the quick sort option or a field well.

To quickly sort dimensions and measures

- Do one of the following:
 - Choose the sort icon that appears near the field name on either axis. In direct queries, this icon appears for any data type. For SPICE, this icon is available only for datetime, numeric, and decimal data types.

- Choose the field name and then choose the sort option from the menu. If the label doesn't display on the axis, check the visual format to see if the axis is set to display labels. The display labels are automatically hidden on smaller visuals. You might need to make the visual large enough to display labels.

To sort by using an off-visual metric

1. Open the analysis with the visual that you want to sort. Visuals pane will be open by default.
2. Choose a field well that supports sorting, then choose **Sort by, Sort options**.
3. On the **Sort options** pane, sort by specific fields, choose an aggregation, or sort ascending or descending, or do a combination of these.
4. Choose **Apply** to save your changes. Or choose **Clear** to start over or **Cancel** to go back.

To sort by using a field well

1. Open the analysis with the visual that you want to sort. Visuals pane will be open by default.
2. Choose a field well that supports sorting.
3. On the field well menu, choose **Sort**, and then choose the ascending or descending sort order icon.

Using themes in Amazon Quick Sight

In Amazon Quick Sight, a *theme* is a collection of settings that you can apply to multiple analyses and dashboards. Amazon Quick Sight includes some themes, and you can add your own by using the theme editor. You can share themes with permissions levels set to user or owner. Anyone who has access to the theme can apply it to analyses and dashboards, or use **Save as** to make their own copy of it. Theme owners can also edit the theme and share it with others.

An analysis can have only one theme applied. If you apply a theme to an analysis (by using the **Apply** button), it instantly changes it for everyone—both analysis and dashboard viewers. To explore and save color options without applying them, avoid editing and saving the applied theme.

All colors come in pairs of background and foreground colors. The foreground colors are meant to specifically appear above their matching background color, so choose something that contrasts well.

The following table defines the different settings.

Group	Setting	What the setting changes
Main	Primary background	The background color used for visuals and other high emphasis UI.
Main	Primary foreground	The color of text and other foreground elements that appear over the primary background regions such as grid lines, borders, table banding, icons, and so on.
Main	Secondary background	The background color used for the sheet background and sheet controls.
Main	Secondary foreground	The foreground color used for any sheet title, sheet control text, or UI that appears over the secondary background.
Main	Accent	This setting is used as an interactive hint for the following: - Buttons - Borders around the selected visual - Loading indicators - Narration customizations

Group	Setting	What the setting changes
		• Links • Filter panes for embedded dashboards
Main	Accent foreground	The foreground color applies to any text or other elements that appear over the accent color.
Main	Font	The font to use for all of the text. You can choose from a variety of fonts supported by Amazon Quick Sight.
Data	Data colors	These are the data colors that charts rotate through when assigning colors to groups. You can add or remove colors to this list, or choose a color to change it.
Data	Min max gradient	The default minimum and maximum gradient colors to use when a gradient is used as a scale, for example in heat maps.

Group	Setting	What the setting changes
Data	Empty fill color	This is the color used with your data colors to indicate a lack of data. For example, this color appears in the empty portion of the progress bars that are shown in key performance indication (KPI) and gauge charts, or for empty heat map cells.
Layout	Border	This setting toggles the border around the visuals that aren't currently selected. The selected visual's border still displays the accent color.
Layout	Margin	This setting toggles the space between the sheet boundaries and the visuals.
Layout	Gutter	This setting shows or hides the space between visuals in the grid.
Other	Success Success foreground	These colors are used for success messages, for example the check mark for a successful download.
Other	Warning Warning foreground	These colors are used for warning and informational messages.
Other	Danger Danger foreground	These colors are used for error messages.

Group	Setting	What the setting changes
Other	Dimension Dimension foreground	These colors are used for the names of fields that are identified as dimensions. This option also sets the color for dimensions in the filter panel of embedded dashboards.
Other	Measure Measure foreground	These colors are used for the names of fields that are identified as measures. These colors also apply to measures in the filter panel of embedded dashboards.

To take a short tour of the theme viewer and editor

1. Open the [Quick console](#).
2. Open an analysis, or create a new one. You must have an analysis open to work with themes. However, the view you see with the theme applied is only a preview.

Themes are separate from analyses. No changes are made to your analysis, even when you save a theme.

3. Choose **Edit** from the application bar, and then choose **Themes**. The themes panel opens.
4. The list of themes shows the following:
 - **Applied theme** shows the theme that is currently applied to this analysis and its dashboards.
 - **My themes** shows themes that you created and themes that are shared with you.
 - **Starter themes** shows themes created by Amazon Quick Sight.
5. Each theme has context menu that you can access from the ... icon.

The actions that are available to you on each theme depend on your level of access.

- **Theme owners** – If you created the theme, or someone shared it with you and made you an owner, you can do the following:

- **Edit** – Change the settings for the theme, and save them.
 - **Save** – Save changes you made to the theme. If you edit the applied theme save your changes, the new theme settings apply to all the analyses and dashboards that use it. An informational message displays before you overwrite an applied theme.
 - **Share** – Share the theme and assign user or owner permissions to other people.
 - **Delete** – Delete a theme. You can't undo this action. An informational message displays before you confirm deletion.
 - **Theme users** – If someone shared the theme with you, or if it's an Amazon Quick Sight theme, you can do the following:
 - **Apply** – Apply the theme to the current analysis. This option also applies the theme to dashboards created from the analysis. An informational message displays before you overwrite an applied theme.
 - **Save as** – Save the current theme to another name, so you can edit it.
 - **Analysis authors** – If you have access to the analysis, but not the theme, you can do the following:
 - You can see the analysis with the theme applied.
 - You can see the theme in the **Theme** panel.
 - You can use **Save as** to create your own copy of the theme.
 - **Dashboard viewers** – If you have access to the dashboard, but not the theme, you can do the following:
 - You can see the dashboard with the theme applied.
 - You can't see the theme or its settings. Dashboard users can't see the **Theme** panel.
6. To explore a theme's settings, choose the icons on the left to see settings for colors.

The following procedure walks you through creating a theme. You can start on the analysis, or a copy of the analysis, that you want to use to preview the colors. Or you can start a new analysis. After you save the theme, you can apply it to the current analysis or to other analyses. If you share it, other people can use it too.

To use the theme editor

1. Open the [Quick console](#).
2. Open an analysis, or create a new one. Choose **Edit** from the application bar, and then choose **Themes**. The **Themes** panel opens.

You must have an analysis open to work with themes. However, the view you see with the theme applied is only a preview. Themes are separate from analyses. No changes are made to your analysis, even when you save a theme.

3. Choose **Main**. The color picker used in each of these settings is the standard one used throughout Amazon Quick Sight.

Set colors for **Primary background** and **Primary foreground** to use in visuals and other high impact UI.

Set colors for **Secondary background** and **Secondary foreground** to use in sheets and sheet controls.

Set colors for **Accent** and **Accent foreground** to use in interactive hints including buttons, borders around selected visuals, loading indicators, narration customizations, links, and the filter pane in embedded dashboards.

4. Choose **Data**.

Set the **Colors** to use as data colors. Charts rotate through these when assigning colors. You can add or delete colors, or change the order they're in by dragging and dropping. To change an existing color, select it to open the color editor.

Set colors for **Min max gradient** to use when a gradient is used as a scale, for example in heat maps.

Set the color for **Empty fill** to use when showing a lack of data, for example the unfilled part of a progress bar.

5. Choose **Layout**.

Enable or disable the **Border** check box to show or hide the border around the visuals that aren't currently selected.

Enable or disable the **Margin** check box to show or hide the space between the sheet boundaries and the visuals.

Enable or disable the **Gutter** check box to show or hide the space between visuals in the grid.

6. Choose **Other**.

Set the color for **Success** to use in success messages, for example when you successfully download a .csv file. The success foreground color isn't currently used.

Set the color for **Warning** to use in warning and informational messages. The warning foreground color isn't currently used.

Set the color for **Danger** to use in error messages. The danger foreground color isn't currently used.

Set the color for **Dimension** to use for the names of fields that are identified as dimensions. This option also sets the color for dimensions in the filter panel of embedded dashboards.

Set the color for **Measure** to use for the names of fields that are identified as measures. This option also sets the color for measures in the filter panel of embedded dashboards.

7. To save the theme, choose **Main** and give the new theme a name, and then choose **Save** at the upper-right of the browser.


Saving a theme doesn't apply it to the analysis, even though you can see a preview of the colors that uses the current analysis.

8. To share the theme, save or close the theme you are viewing. Find the theme in your theme collection. Choose **Share** from the context menu (...).
9. To apply the theme, save or close the theme you are viewing. Find the theme in your theme collection. Choose **Apply** from the context menu (...).

Accessing Amazon Quick Sight using keyboard shortcuts


You can use the following keyboard shortcuts to navigate a Amazon Quick Sight dashboard or analysis:

- Use the TAB key to navigate among menu options or visuals.
- Use the Shift+TAB keys to move backward to the previous selection.
- Use the Enter key to select a visual or menu option.
- Use the ESC key to clear the selection from a visual or menu item.




Using shortcuts within a visual

You can use the TAB, Shift+TAB, and Enter keys to navigate and select different fields within a selected visual. For example, say that you want to use a link that's a part of your visuals title. To do this, select the visual that you want, then use the TAB key until just the link is selected. Then, use the Enter key to click on the link.



You can also use these keyboard shortcuts to navigate and enter the on-visual menu on the upper-right corner of a visual. To do this, select the visual that you want and use the TAB key to get to the field that you want to select. If you miss the field that you want, use the Shift+TAB keys to go back a field.



Sharing and subscribing to data in Amazon Quick Sight with dashboards and reports

A *dashboard* is a read-only snapshot of an analysis that you can share with other Amazon Quick Sight users for reporting purposes. A dashboard preserves the configuration of the analysis at the time you publish it, including such things as filtering, parameters, controls, and sort order. The data used for the analysis isn't captured as part of the dashboard. When you view the dashboard, it reflects the current data in the data sets used by the analysis.

When you share a dashboard, you specify which users have access to it. Users who are dashboard viewers can view and filter the dashboard data. Any selections to filters, controls, or sorting that users apply while viewing the dashboard exist only while the user is viewing the dashboard, and aren't saved after it's closed. Users who are dashboard owners can edit and share the dashboard, and optionally can edit and share the analysis. If you want them to also edit and share the data set, you can set that up in the analysis.

A shared dashboard can also be embedded in a website or app, if you are using Enterprise edition. For more information about embedded dashboards, see [Embedded analytics for Amazon Quick Sight](#).

Use the following sections to learn how to publish and share dashboards, subscribe to threshold alerts, and send and subscribe to dashboard email reports.

Topics

- [Publishing dashboards](#)
- [Sharing Amazon Quick Sight dashboards](#)
- [Using Quick connectors in dashboard visuals](#)
- [Sharing your view of a Amazon Quick Sight dashboard](#)
- [Scheduling and sending Quick Sight reports by email](#)
- [Subscribing to email reports in Amazon Quick Sight](#)
- [Working with threshold alerts in Amazon Quick Sight](#)
- [Printing a dashboard or analysis](#)
- [Exporting Amazon Quick Sight analyses or dashboards as PDFs](#)
- [Error codes for failed PDF export jobs](#)
- [Organizing assets into folders for Amazon Quick](#)

Publishing dashboards

When you publish an analysis, that analysis becomes a dashboard that can be shared and interacted with by users of your Amazon Quick account or, in some cases, with anonymous users that aren't on your account. You can choose to publish one sheet of an analysis, all sheets in the analysis, or any other combination of sheets that you want. When you publish an interactive sheet, that sheet becomes an interactive dashboard that users can interact with. When you publish a pixel perfect report sheet, the sheet becomes a pixel perfect report that generates and saves a snapshot of the report's data when you schedule a report in Amazon Quick Sight. You can publish a dashboard that contains any combination of interactive sheets and pixel perfect reports from the same analysis.

For more information about scheduling a report, see [Scheduling and sending Quick Sight reports by email](#).

For more information about viewing a report's snapshots, see [Consuming pixel perfect reports in Amazon Quick Sight](#).

Use the following procedure to publish and optionally share a dashboard. You can also use this procedure to rename a published dashboard. A renamed dashboard retains its security and emailed report settings.

1. Open the analysis that you want to use. Choose **Publish**.
2. Do one of the following:
 - To create a new dashboard, choose **New dashboard**, and then type a dashboard name.
 - To replace an existing dashboard, do one of the following. Replacing a dashboard updates it without altering security or emailed report settings.
 - To update it with your changes, choose **Replace an existing dashboard** and then choose a dashboard from the list.
 - To rename it, choose **Replace an existing dashboard**, choose a dashboard from the list, and then select the pencil icon. Enter a new name to rename the existing dashboard and click the checkmark or press enter to confirm. When you publish a dashboard after renaming, it also saves any changes you made to the analysis. Changes to the analysis or dashboard are not persisted until you **Publish**. An initial version of a dashboard must be published in order to rename it.
3. (Optional) Choose the sheets that you want to publish in the **SHEETS** dropdown. When you select sheets to add to the new dashboard, the dropdown shows how many sheets are selected for publishing. The default option is **ALL SHEETS SELECTED**.

If you are replacing an existing dashboard, the sheets that are already published to the existing dashboard are pre-selected in the dropdown, unless you are publishing from an analysis you have not previously published from. You can make changes to this by selecting or de-selecting sheets from the dropdown list.

4. (Optional) Add comments on the changes you have made in the notes section, which is available to view under [Version History](#).
5. (Optional) To allow dashboard readers to share data stories, choose **Allow sharing data stories**. For more information about data stories, see [Working with data stories in Amazon Quick Sight](#).
6. (Optional) Open **More Settings**. These options are only available if at least one sheet in the new dashboard is an interactive sheet.

 Note

This is a scrollable window. Scroll down in the **Publish a dashboard** window to view all available options.

There are some options that you can turn off to simplify the experience for this dashboard, as follows:

- For **Dashboard options**:
 - Leave **Expand on-sheet controls by default** cleared to show a simplified view. This is disabled by default. To show the controls by default, turn on this option.
 - Clear **Enable advanced filtering on the left pane** to remove the ability for dashboard viewers to filter the data themselves. If they create their own filters, the filters exist only while the user is viewing the dashboard. Filters can't be saved or reused.
 - Clear **Enable on-hover tooltip** to turn off tooltips.
- For **Visual options**:
 - Clear **Enable visual menu**, to turn off the on-visual menu entirely.
 - Clear **Enable download options** if your dashboard viewers don't need to be able to download data from the visuals in the dashboard. The CSV file includes only what is currently visible in the visual at the time they download it. The viewer downloads data by using the on-visual menu on each individual visual.
 - Clear **Enable maximize visual option** to turn off the ability to enlarge visuals to fill the screen.
- For **Data point options**:
 - Clear **Enable drill up/down** if your dashboard doesn't offer drillable field hierarchies.
 - Clear **Enable on-click tooltip** to turn off tooltips that appear when the reader chooses (clicks on) a data point.
 - Clear **Enable sort options** to turn off sorting controls.

7. Choose **Publish dashboard**.

If you renamed the existing dashboard, the top of the screen refreshes to show the new name.

8. (Optional) Do one of the following:

- To publish a dashboard without sharing, choose **x** at the upper right of the **Share dashboard with users** screen when it appears. You can always share the dashboard later by choosing **File>Share** from the application bar.
- To share the dashboard, follow the procedure in [Sharing Amazon Quick Sight dashboards](#).

After you complete these steps, you complete creating and sharing the dashboard. Subscribers of the dashboard receive email that contains a link to the dashboard. Groups don't receive invitation emails.

Copying an Amazon Quick Sight dashboard

If you have co-owner access or **Save as** privileges on an existing dashboard, you can copy it. To do this, create a new analysis from the dashboard and then create a new dashboard from the analysis that you copied.

After you save the original dashboard as a new analysis, you can collaborate on it by sharing the new analysis with other users. For example, you can use this workflow to preserve a production version of the dashboard, while also developing or testing a new version of it.

To copy a dashboard

1. Sign in to Quick and choose **Dashboards** from the homepage.
2. Open the dashboard that you want to duplicate.
3. At upper right, choose **Save As**, and then enter a name for the new analysis. When you save an existing dashboard using **Save As**, it creates an analysis based on the dashboard.

Note

If you can't see **Save as**, check with your administrator that you have the right permissions.

4. (Optional) Make changes to the new analysis.
5. (Optional) Share the analysis with other users so you can collaborate on changes. All users who have access can make changes to the new analysis.

To share the analysis with other users, choose **Share** from the top right corner of the page, and then choose **Share analysis**.

6. (Optional) Create a new dashboard with your changes to the new analysis by choosing **Share**, and then choosing **Publish Dashboard**.

For more information, see the following:

- [Sharing Amazon Quick Sight dashboards](#)
- [Sharing Quick Sight analyses](#)

Deleting an Amazon Quick Sight dashboard

When you delete an Amazon Quick Sight dashboard, the dashboard is permanently removed from your account and all folders that the dashboard was a part of. You will no longer be able to access the deleted dashboard. You can only delete dashboards that you own or co-own. Use the following procedure to delete a dashboard.

To delete a dashboard

1. On the **Dashboards** tab of the Amazon Quick homepage, choose the details icon (vertical dots :) on the dashboard that you want to delete.
2. Choose **Delete**. Then choose **Delete** again to confirm that you want to delete the dashboard. Deleting a dashboard permanently deletes the dashboard from your account, and the dashboard will disappear from all folders that it belonged to. You can still access and create other dashboards from the analysis that the deleted dashboard was published from.

Publishing a previous version of an Amazon Quick Sight dashboard

Each time you make updates to an analysis and publish it, a new version of the Amazon Quick Sight dashboard is created. To revert back to a previous version of a dashboard, you can search for it under the dashboard's **Version History** and publish the former version you are interested in. Each dashboard can store up to 1000 versions that are never deleted. Use the following procedure to publish a previous version of a dashboard.

To publish a previous version of a dashboard

1. On the **Dashboards** tab of the Amazon Quick homepage, choose the dashboard that you want to manage.

2. Choose **Version History** on the toolbar on the right. The version of the dashboard that is currently published, as well as previous available versions, will appear in a list. Any comments that were added in the notes section will appear with the respective version.
3. Select the version of the dashboard you are interested in. You can see when this version was published and which user published it.
4. To revert to this version, select **Publish**. Click **Confirm** to publish the version.

Sharing Amazon Quick Sight dashboards

By default, dashboards in Amazon Quick Sight aren't shared with anyone and are only accessible to the owner. However, after you publish a dashboard, you can share it with other users or groups in your Amazon Quick account. You can also choose to share the dashboard with everyone in your Quick account and make the dashboard visible on the Quick homepage for all users in your account. Additionally, you can copy a link to the dashboard to share with others who have access to it.

Important

Users who have access to the dashboard can also see the data used in the associated analysis.

After you share a dashboard, you can review the other users or groups that have access to it and control the type of access they have. You can revoke access to the dashboard for any user. You can also remove yourself from it.

You can also embed interactive dashboards and visuals in websites and apps by copying the dashboard or visual embed code and pasting it in your application. For more information, see [Embedding Amazon Quick Sight visuals and dashboards for registered users with a 1-click embed code](#).

Granting access to a dashboard

You can share dashboards and visuals with specific users or groups in your account or with everyone in your Amazon Quick account. Or you share them with anyone on the internet. You can share dashboards and visuals by using the Quick console or the Quick Sight API. Access to a shared visual depends on the sharing settings that are configured for the dashboard that the visual

belongs to. To share and embed visuals to your website or application, adjust the sharing settings of the dashboard that it belongs to. For more information, see the following:

- [Granting individual Amazon Quick Sight users and groups access to a dashboard in Amazon Quick Sight](#)
- [Granting everyone in your Amazon Quick Sight account access to a dashboard](#)
- [Granting anyone on the internet access to an Amazon Quick Sight dashboard](#)
- [Granting everyone in your Amazon Quick account access to a dashboard with the Quick Sight API](#)
- [Granting anyone on the internet access to an Amazon Quick Sight dashboard using the Quick Sight API](#)

Granting individual Amazon Quick Sight users and groups access to a dashboard in Amazon Quick Sight

Use the following procedure to grant access to a dashboard.

To grant users or groups access to a dashboard

1. Open the published dashboard and choose **Share** at upper right. Then choose **Share dashboard**.
2. In the **Share dashboard** page that opens, do the following:
 - a. For **Invite users and groups to dashboard** at left, enter a user email or group name in the search box.

Any users or groups that match your query appear in a list below the search box. Only active users and groups appear in the list.

- b. For the user or group that you want to grant access to the dashboard, choose **Add**. Then choose the level of permissions that you want them to have.

You can select **Viewer** or **Co-owner**, depending on the user's Quick role. The available permissions for each role are as follows:

- **Readers** – Quick readers can only be granted **Viewer** access to dashboards. They can view, export, and print the dashboard, but they can't save the dashboard as an analysis. They can view, filter, and sort the dashboard data. They can also use any controls or custom actions that are on the dashboard. Any changes that they make to the

dashboard exist only while they are viewing it, and aren't saved after they close the dashboard.

- **Authors** – Quick authors can be granted **Viewer** or **Co-owner** access to dashboards.
 - Authors with Viewer access can view, export, and print the dashboard. They can view, filter, and sort the dashboard data. They can also use any controls or custom actions that are on the dashboard. Any changes that they make to the dashboard exist only while they are viewing it, and aren't saved after they close the dashboard.

However, they can save the dashboard as an analysis, unless the dashboard owner specifies otherwise. This privilege grants them read-only access to the datasets so that they can create new analyses from them. The owner has the option to provide them with the same permissions to the analysis. If the owner wants them also to edit and share the datasets, the owner can set that up inside the analysis.

- Authors with Co-owner access can view, export, and print the dashboard. They can also edit, share, and delete it. They can also save the dashboard as an analysis, unless the dashboard owner specifies otherwise. This privilege grants them read-only access to the datasets so that they can create new analyses from them. The owner has the option to provide them with the same permissions to the analysis. If the owner wants them to also edit and share the datasets, the owner can set that up inside the analysis.
- **Groups** – Quick groups can only be granted **Viewer** access to dashboards. They can view, export, and print the dashboard, but they can't save the dashboard as an analysis.

After you add a user or group to the dashboard, you can see information about them in the **Manage permissions** section, under **Users & Groups**. You can see their user name, email, permission level, and "save as" privileges.

To allow a user or group to save the dashboard as an analysis, turn on **Allow "save as"** in the **Save as Analysis** column.

- c. To add more users to the dashboard, enter another user email or group name in the search box and repeat steps A and B.

Granting everyone in your Amazon Quick Sight account access to a dashboard

Alternatively, you can share your Amazon Quick Sight dashboard with everyone in your account. When you do this, everyone in your account can access the dashboard, even if they weren't granted

access individually and assigned permissions. They can access the dashboard if they have a link to it (shared by you) or if it's embedded.

Sharing the dashboard with everyone in your account doesn't affect email reports. For example, suppose that you choose to share the dashboard with everyone in your account. Suppose also that you choose **Send email report to all users with access to dashboard** when setting up an email report for the same dashboard. In this case, the email report is sent only to people who have access to the dashboard. They receive access either through someone explicitly sharing it with them, through groups, or through shared folders.

To grant everyone in your account access to a dashboard

1. Open the published dashboard and choose **Share** at upper right. Then choose **Share dashboard**.
2. In the **Share dashboard** page that opens, for **Enable access for** at bottom left, toggle on **Everyone in this account**. Accounts that sign in with an Active Directory can't access the **Everyone in this account** switch. Accounts that use Active Directory can enable this setting with an `UpdateDashboardPermissions` API call. For more information on `UpdateDashboardPermissions`, see [UpdateDashboardPermissions](#) in the *Amazon Quick Sight API Reference*.
3. (Optional) Toggle on **Discoverable in Quick Sight**.

When you share a dashboard with everyone in the account, owners can also choose to make the dashboard discoverable in Quick Sight. A dashboard that's discoverable appears in everyone's list of dashboards on the **Dashboards** page. When this option is turned on, everyone in the account can see and search for the dashboard. When this option is turned off, they can only access the dashboard if they have a link or if it's embedded. The dashboard doesn't appear on the **Dashboards** page, and users can't search for it.

Granting anyone on the internet access to an Amazon Quick Sight dashboard

Applies to: Enterprise Edition

You can also share your Amazon Quick Sight dashboard with anyone on the internet from the **Share** menu in the Amazon Quick console. When you do this, anyone on the internet will be able to

access the dashboard, even if they aren't a registered user on your Quick account, when you share the dashboard link or embed the dashboard.

Use the following sections to grant anyone on the internet access to dashboard when you share it.

Topics

- [Before you start](#)
- [Granting anyone on the internet access to a dashboard](#)
- [Updating a publicly shared dashboard](#)
- [Turning off public sharing settings](#)

Before you start

Before you can share a dashboard with anyone on the internet, make sure to do the following:

1. Turn on session capacity pricing on your account. If you have not turned on session capacity pricing on your account, you won't be able to update your account's public sharing settings.
2. Assign public sharing permissions to an administrative user in the IAM console. You can add these permissions with a new policy or you can add the new permissions to an existing user.

The following sample policy provides permissions for use with `UpdatePublicSharingSettings`.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": "quicksight:UpdatePublicSharingSettings",
      "Resource": "*",
      "Effect": "Allow"
    }
  ]
}
```

Accounts that don't want users with administrator access to use this feature can add an IAM policy that denies public sharing permissions. The following sample policy denies permissions for use with `UpdatePublicSharingSettings`.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": "quicksight:UpdatePublicSharingSettings",
      "Resource": "*",
      "Effect": "Deny"
    }
  ]
}
```

For more information on using IAM with Quick Sight, see [Using Quick with IAM](#).

You can also use the "Deny" policy as a Service Control Policy (SCP) if you don't want any of the accounts in your organization to have the public sharing feature. For more information, see [Service control policies \(SCPs\)](#) in the *AWS Organizations User Guide*.

3. Turn on public sharing on your Amazon Quick account.
 1. From the Amazon Quick start page, choose your user icon at the upper right of your browser window, and then choose **Manage Quick**.
 2. In the page that opens, scroll down to the **Permissions** section.
 3. Choose **Public access to dashboards** at left.
 4. On the page that opens, choose **Anyone on the internet**.

When you turn on this setting, a pop up will appear asking you to confirm your choice. Once you've confirmed your choice, you can grant the public access to specific dashboards and share those dashboards with them with a link or by embedding the dashboard in a public application, wiki, or portal.

Granting anyone on the internet access to a dashboard

To grant anyone on the internet access to a dashboard

1. In Quick, open the published dashboard that you want to share. You must be the owner or a co-owner of the dashboard.

2. In the published dashboard, choose the **Share** icon at upper-right, and then choose **Share dashboard**.
3. In the **Share dashboard** page that opens, choose **Anyone on the internet (public)** in the **Enable access for** section at bottom-left.

This setting allows you to share the dashboard with anyone on the internet with the share link or when embedded. Turning on this switch also automatically turns on the **Everyone in this account** option, which means that the dashboard will be shared with anyone in your Quick account. If you do not want this, turn off this option.

4. In the **Allow public access** pop-up that appears, enter `confirm` in the box to confirm your choice, and then choose **Confirm**.

After you confirm your dashboard's access settings, an orange **PUBLIC** tag appears at upper right of your dashboard in the Amazon Quick console. Additionally, an eye icon appears on the dashboard on the Quick Sight Dashboards page, both in tile and list view.

Note that when public access is turned on, the dashboard can only be accessed using the link or when embedded using the embed code. For more information about sharing a link to the dashboard, see [Sharing a link a shared dashboard](#). For more information about embedding dashboards for anyone on the internet, see [Embedding Amazon Quick Sight visuals and dashboards for anonymous users with a 1-click embed code](#).

Updating a publicly shared dashboard

Use the following procedure to update a shared dashboard that can be accessed by anyone on the internet.

To update a public dashboard:

1. From the Amazon Quick start page, choose the analysis that is tied to the dashboard that you want to update and make your desired changes. You must be the owner or a co-owner of the analysis.
2. In the analysis, choose **Publish**.
3. In the pop-up that appears, choose **Replace an existing dashboard** and select the public dashboard that you want to update.
4. To confirm your choice, enter `confirm` and then choose **Publish dashboard**.

Once you choose **Publish dashboard**, your public dashboard is updated to reflect the new changes.

Turning off public sharing settings

You can turn off public sharing settings for dashboards at anytime. You can turn off public sharing for an individual dashboard, or for all dashboards in your account. Visual sharing settings are determined at the dashboard level. If you turn off public sharing settings to a dashboard that holds a visual that you are embedding, users won't be able to access the visual.

The following table describes the different scenarios for when a dashboard is publicly available.

Account-level public setting	Dashboard-level public setting	Public access	Visual indicators
Off	Off	Off	None
On	Off	Off	None
On	On	Yes	An orange badge appears on the dashboard and an eye icon appears on the dashboard in the Dashboards page.
Off	On	No	A grey badge appears on the dashboard and an eye icon with a slash appears on the dashboard in the Dashboards page. It can take up to two minutes for

Account-level public setting	Dashboard-level public setting	Public access	Visual indicators
			a dashboard's public access to be revoked.

To turn off public sharing for a single dashboard

1. In Amazon Quick, open the published dashboard that you no longer want to share. You must be the owner or a co-owner of the dashboard.
2. In the published dashboard, choose the **Share** icon at upper-right, and then choose **Share dashboard**.
3. In the **Share dashboard** page that opens, toggle off the **Anyone on the internet (public)** switch in the **Enable access for** section at bottom-left.

This action will remove public access to the dashboard. It will now only be accessible to users that it has been shared with.

To turn off public sharing settings for all dashboards in a Quick user account

1. From the Amazon Quick start page, choose your user icon at upper right of your browser window, and then choose **Manage Quick**.
2. In the page that opens, scroll down to the **Permissions** section.
3. Choose **Public access to dashboards** at left.
4. On the page that opens, toggle off the **Anyone on the internet** switch.

When you disable public sharing settings from the **Public sharing** menu, a pop-up will appear asking you to confirm your choice. Select **I have read and acknowledge this change** and then choose **Confirm** to confirm your choice.

This action will remove public access to every dashboard on your account. Dashboards that were visible to anyone on the internet will now only be accessible to users that each dashboard has been shared with. Individual dashboards that have their public settings turned on will have a gray badge and the eye icon that appears on the **Dashboards** page will have a strike through

it to indicate that the account level public settings are disabled and that the dashboard can't be viewed. It can take up to two minutes for a dashboard's public access to be revoked.

If your session capacity pricing subscription has expired, public sharing settings will be automatically removed across your account. Renew your subscription to restore access to public sharing settings.

Granting everyone in your Amazon Quick account access to a dashboard with the Quick Sight API

Intended audience: Amazon Quick developers

Alternatively, you can grant everyone in your account access to the dashboard with the Quick Sight API using the `UpdateDashboardPermissions` operation.

The following example API request illustrates how to do so using an AWS CLI command. It grants link permissions on the dashboard in your account, and allows the following operations: `DescribeDashboard`, `QueryDashboard` and `ListDashboard`.

```
aws quicksight update-dashboard-permissions \  
--aws-account-id account-id \  
--region aws-directory-region \  
--dashboard-id dashboard-id \  
--grant-link-permissions  
Principal="arn:aws:quicksight:aws-directory-region:account-id:namespace/default",  
Actions="quicksight:DescribeDashboard, quicksight:QueryDashboard,  
quicksight:ListDashboardVersions"
```

The response for the preceding request looks similar to the following.

```
{  
  "Status": 200,  
  "DashboardArn": "arn:aws:quicksight:AWSDIRECTORYREGION:ACCOUNTID:dashboard/  
DASHBOARDID",  
  "DashboardId": "DASHBOARDID",  
  "LinkSharingConfiguration": {  
    "Permissions": [  
      {  
        "Actions": [  

```

```

    "quicksight:DescribeDashboard",
    "quicksight:ListDashboardVersions",
    "quicksight:QueryDashboard"
  ],
  "Principal": "arn:aws:quicksight:AWSDIRECTORYREGION:ACCOUNTID:namespace/default"
}
],
"Permissions": [
  // other dashboard permissions here
],
"RequestId": "REQUESTID"
}

```

You can also prevent all users in your account from accessing the dashboard using the same API operation. The following example request illustrates how by using a CLI command.

```

aws quicksight update-dashboard-permissions \
--aws-account-id account-id \
--region aws-directory-region \
--dashboard-id dashboard-id \
--revoke-link-permissions
Principal="arn:aws:quicksight:aws-directory-region:account-id:namespace/default",
Actions="quicksight:DescribeDashboard, quicksight:QueryDashboard,
quicksight:ListDashboardVersions"

```

For more information, see [UpdateDashboardPermissions](#) in the *Amazon Quick API Reference*.

When all users in a Quick user account are granted access to the dashboard, the following snippet is added to AWS CloudTrail log as part of the eventName `UpdateDashboardAccess`, and the eventCategory `Management`.

```

"linkPermissionPolicies":
[
{
  "principal": "arn:aws:quicksight:AWSDIRECTORYREGION:ACCOUNTID:
    namespace/default",
  "actions":
  [
    "quicksight:DescribeDashboard",
    "quicksight:ListDashboardVersions",
    "quicksight:QueryDashboard"
  ]
}
]

```

```
]
}
]
```

Granting anyone on the internet access to an Amazon Quick Sight dashboard using the Quick Sight API

Alternatively, you can grant anyone on the internet access to the dashboard with the Amazon Quick Sight API using the `UpdateDashboardPermissions` operation.

Before you begin, make sure to grant everyone in your account access to the dashboard. For more information, see [Granting everyone in your Amazon Quick account access to a dashboard with the Quick Sight API](#).

The following example API request illustrates how to grant anyone on the internet access to a dashboard using an AWS CLI command. It grants link permissions on the dashboard in your account, and allows the following operations: `DescribeDashboard`, `QueryDashboard` and `ListDashboardVersions`.

```
aws quicksight update-dashboard-permissions
--aws-account-id account-id
--region aws-directory-region
--dashboard-id dashboard-id
--grant-link-permissions
Principal="arn:aws:quicksight::publicAnonymousUser/*",
Actions="quicksight:DescribeDashboard, quicksight:QueryDashboard,
quicksight:ListDashboardVersions"
```

The response for the preceding request looks similar to the following.

```
{
  "Status": 200,
  "DashboardArn": "arn:aws:quicksight:AWSDIRECTORYREGION:ACCOUNTID:dashboard/DASHBOARDID",
  "DashboardId": "DASHBOARDID",
  "LinkSharingConfiguration": {
    "Permissions": [
      {
        "Actions": [
          "quicksight:DescribeDashboard",
          "quicksight:ListDashboardVersions",
          "quicksight:QueryDashboard"
        ]
      }
    ]
  }
}
```

```

        ],
        "Principal":
"arn:aws:quicksight:AWSDIRECTORYREGION:ACCOUNTID:namespace/default"
    },
    "Principal": "arn:aws:quicksight:::publicAnonymousUser/*",
    "Actions": [
        "quicksight:DescribeDashboard",
        "quicksight:ListDashboardVersions",
        "quicksight:QueryDashboard"
    ]
    }
    ]
},
"Permissions": [
    // other dashboard permissions here
],
"RequestId": "REQUESTID"
}

```

You can also prevent anyone on the internet from accessing the dashboard using the same API operation. The following example request illustrates how by using a CLI command.

```

aws quicksight update-dashboard-permissions \
--aws-account-id account-id \
--region aws-directory-region \
--dashboard-id dashboard-id \
--revoke-link-permissions
Principal="arn:aws:quicksight:::publicAnonymousUser/*",
Actions="quicksight:DescribeDashboard, quicksight:QueryDashboard,
quicksight:ListDashboardVersions"

```

For more information, see [UpdateDashboardPermissions](#) in the *Amazon Quick API Reference*.

When anyone on the internet is granted access to the dashboard, the following snippet is added to AWS CloudTrail log as part of the eventName `UpdateDashboardAccess`, and the eventCategory `Management`.

```

"linkPermissionPolicies":
[
{
    "principal": "arn:aws:quicksight:::publicAnonymousUser/*",
    "actions":

```

```
[
  "quicksight:DescribeDashboard",
  "quicksight:ListDashboardVersions",
  "quicksight:QueryDashboard"
]
}
```

Sharing a link a shared dashboard

After you grant users access to a dashboard, you can copy a link to it and send it to them. Anyone with access to the dashboard can access the link and see the dashboard.

To send users a link to the dashboard

1. Open the published dashboard and choose **Share** at upper right. Then choose **Share dashboard**.
2. In the **Share dashboard** page that opens, choose **Copy link** at upper left.

The link to the dashboard is copied to your clipboard. It's similar to the following,

```
https://quicksight.aws.amazon.com/sn/accounts/accountid/
dashboards/dashboardid?directory_alias=account_directory_alias
```

Users and groups (or all users on your Quick account) who have access to this dashboard can access it by using the link. If they are accessing Quick for the first time, they will be asked to sign in with their email address or Quick user name and password for the account. After they sign in, they will have access to the dashboard.

View who has access to a shared dashboard

Use the following procedure to see which users or groups have access to the dashboard.

1. Open the published dashboard and choose **Share** at upper right. Then choose **Share dashboard**.
2. In the **Share dashboard** page that opens, under **Manage permissions**, review the users and groups, and their roles and settings.

You can search to locate a specific user or group by entering their name or any part of their name in the search box at upper right. Searching is case-sensitive, and wildcards aren't supported. Delete the search term to return the view to all users.

Revoke access to a shared dashboard

Use the following procedure to revoke user access to a dashboard.

To revoke user access to a dashboard

1. Open the dashboard and choose **Share** at top right. Then choose **Share dashboard**.
2. In the **Share dashboard** page that opens, under **Manage permissions**, locate the user that you want to remove and choose the delete icon at far right.

Using Quick connectors in dashboard visuals

Prerequisites

Before you begin, make sure to [create at least one connector](#).

The connector must meet these requirements:

- Uses the **User Auth** authentication method
- Uses one of the following integrations:
 - Atlassian Jira Cloud
 - Microsoft Outlook
 - Microsoft Teams
 - Salesforce
 - ServiceNow
 - Slack

Enable Quick actions on a dashboard to use connectors

To enable Quick actions on a dashboard to use connectors

1. If a dashboard exists, go to the source analysis of the dashboard. Otherwise, [create a new analysis](#).
2. Choose **Publish**.
3. Choose between **New Dashboard** or **Replace existing dashboard**.
4. Choose the **Enable Quick actions** checkbox under **Dashboard options**.
5. Choose **Publish dashboard**.

Use connectors on a visual

To use connectors on a visual

1. Open a dashboard with the **Enable Quick actions** publishing option turned on.
2. Hover over a visual.
3. Choose the lightening bolt icon.
4. A menu appears with a list of all supported connectors and actions.
5. Choose the desired action from the list.
6. If you have not used the connector before, or if your previous login credentials have expired, an authentication modal will appear. Log in with appropriate credentials for your organization.
7. An **Action** form appears in the right pane.
8. Enter all the information you need to include with the action.
9. Some fields allow the inclusion of autofill values. Choose **Autofill** to open the menu. Choose the values you need and they will be added to your entered text.
 - **Today's date**: Injects today's date
 - **Visual name**: Injects visual name
 - **All**: Injects both of the above
10. Some actions support the ability to include an attachment. You can optionally attach an image of the visual with these actions by selecting the **Visual image** checkbox.
11. Select the action button at the bottom of the form to invoke the action.

Security and customizations

Custom Permissions/Capability Customization

- **Actions** capability: You cannot see or use actions if your user or role is restricted the permission to use the Actions capability

To learn more about custom permissions, see [Custom permissions in Amazon Quick](#).

Row Level Security (RLS) / Column Level Security (CLS)

- You cannot see or use actions on visuals that are based on datasets that use RLS or CLS.

To learn more about RLS, see [Using row-level security in Amazon Quick](#).

To learn more about CLS, see [Using column-level security to restrict access to a dataset](#).

Dashboard publishing options

- Enable Quick actions
 - You cannot see or use actions on any visuals of a dashboard that was published with the **Enable Quick actions** publishing option disabled.

To learn more about dashboard publishing options, see [Publishing dashboards](#).

Limitations

Visual Image attachment support

The following visual types do not support image attachments:

- High charts (when HTML is used)
- ML Insights (when HTML is used)
- Textbox and insights (when HTML is used)
- Custom content

 Note

For these visuals, the **Visual image** checkbox will not appear on the UI.

Sharing your view of a Amazon Quick Sight dashboard

While interacting with a published dashboard, you can choose to share a unique link to the dashboard with only your changes. For example, if you filter the data in the dashboard, you can share what you see with others who have permissions to see the dashboard. That way, they can see what you see, without your having to create a new dashboard.

When others access your view of the dashboard by using the link that you sent them, they see the dashboard exactly as it was when the link was created. They see any parameters, filters, or controls that you changed.

To share your view of a dashboard

1. Open the published dashboard, and make any changes that you want.
2. Choose **Share** at upper right, and then choose **Share this view**.
3. On the **Share using a link** page that opens, choose **Copy link**.
4. Paste the link in an email or IM message to share it with others.

Only people with permissions to see the dashboard in Quick Sight can access the link.

Scheduling and sending Quick Sight reports by email

 Important

Amazon Quick Sight in the Europe (Spain) (eu-south-2) region uses an internal email service (Amazon SES) in the Europe (Ireland) (eu-west-1) to send emails to Quick Sight users. Customer data that's included in scheduled reports, alerts, and other features are passed by email from Europe (Spain) to Europe (Ireland) before it reaches Quick Sight users. As a privacy protection measure, the following features that send customer data in emails have been limited or disabled by default.

- File attachments and sheet previews in Scheduled Report emails. The [download link option](#) is the default.

- Emails that use threshold alerts.
- Anomaly detection alerts.

For more information about AWS privacy features, see [Privacy Features of AWS Services](#).

In Enterprise edition, you can send a dashboard in report form either once or on a schedule (daily, weekly, monthly, or yearly). You can email the reports to users or groups who share your Amazon Quick subscription. To receive email reports, the users or group members must meet the following conditions:

- They are part of your Quick subscription.
- You already shared the dashboard with them.
- Amazon Quick Sight can't send scheduled emails to more than 5,000 members.

Amazon Quick Sight generates a custom email snapshot for each user or group based on their data permissions, which are defined in the dashboard. Row Level Security (RLS), Column Level Security (CLS) and Dynamic Default Parameters for email reports works for both scheduled and ad hoc (one-time) emails.

Quick authors can run scheduled reports with the **Report now** button in the Quick console or with the [StartDashboardSnapshotJobSchedule](#) API.

Subscribers who are readers see an option for **Reports** on the dashboard when an email report is available for that dashboard. They can use the **Schedules** menu to subscribe to or unsubscribe from the emails. For more information, see [Subscribing to email reports in Amazon Quick Sight](#).

You can create up to five schedules for each dashboard.

Quick Sight dashboard viewers can also schedule their own reports for themselves from a Quick Sight dashboard. For more information about reader generated reports, see [Creating a reader generated report in Amazon Quick Sight](#).

Use the following topics to learn more about email report settings and report billing.

Topics

- [Configuring email report settings for a Quick Sight dashboard](#)
- [How billing works for email reports](#)

Configuring email report settings for a Quick Sight dashboard

Applies to: Enterprise Edition

In Amazon Quick Enterprise edition, you can email a report from any sheet in a dashboard. You can send reports from interactive dashboards and pixel perfect report sheets. Schedules include settings for when to send them, the contents to include, and who receives the email. You can view a sample report and a list of the datasets used in the report. To set up or change the schedule sent from a dashboard, make sure that you're an owner or co-owner of the dashboard.

If you have access to the dashboard, you can change your subscription options by opening your view of the dashboard. For more information on how this works, see [Subscribing to email reports in Amazon Quick Sight](#).

Scheduling options that are available for an email report include the following:

- **Once (Does not repeat)** – Sends the report only once at the date and time that you choose.
- **Daily** – Repeats daily at the time that you choose.
- **Weekly** – Repeats each week on the same day or days at the time that you choose. You can also use this option to send reports in weekly intervals, such as every other week or every three weeks.
- **Monthly** – Repeats each month on the same day of the month at the time that you choose. You can also use this option to send reports on specific days of the month, such as the second Wednesday or the last Friday of each month.
- **Yearly** – Repeats each year on the same day of the month or months selected at the time that you choose. You can also use this option to send reports on specific days or sets of days in selected months. For example, you can configure a report to be sent on the first Monday of January, March, and September, or on July 14th, or on the second day of February, April, and June each year.
- **Custom** – Configure your own scheduled report that best fits your business needs.

You can customize the title of the report, the optional email subject, and the body text.

Although you can configure the report so that everyone who has access receives a copy, this is not usually the best plan. We recommend limiting automated emails, especially those sent to groups.

You can start with a small number of subscribers by choosing specific people from the access list. Verify your company's policy before subscribing anyone to a subscription.

You can directly add people to a report subscription in these ways:

- (Recommended) Choose recipients from the provided access list to specify and maintain a list of people who you want to email reports to. You can use the search box to find people by email or group name.
- To send reports to all of the dashboard's subscribers, choose **Send email report to all users with access to dashboard** when prompted.

Anyone else who wants to get the emails can open the dashboard and set their own subscription options to either opt in or opt out.

Important

When you share the dashboard with new Quick user names or groups, they automatically start receiving the email reports. If you don't want this to happen, you need to edit the report settings each time you add people to the dashboard.

For existing email schedules, you can pause the schedule in Amazon Quick Sight while you make changes. In the **Schedules** pane, you can pause or resume a scheduled report with the toggle that appears under each report. Pausing a report does not delete the report's schedule from Quick Sight.

If your report includes custom visuals, be aware that you can't include images from a private network in an email report, even if you can access the images. If you want to include an image, use a publicly available one.

Before you begin, make sure that you are using Amazon Quick Enterprise edition and that you have shared the dashboard with intended recipients.

To create or change an email report

1. Open Quick and choose **Dashboards** on the navigation pane at left.
2. Open a dashboard to configure its email report.
3. At top right, choose **Schedules**, and then choose **Schedules**.

4. Choose **ADD SCHEDULE**.
5. In the **New schedule** pane that appears, enter the schedule name. Optionally, add a description for the new schedule.
6. In the **Content** tab, toggle the **PDF**, **CSV**, or **Excel** switches to choose the report format. CSV and Excel format are currently supported for pixel perfect reports.
7. In the **Sheet** dropdown on the **Content** tab, choose the sheet that you want to schedule a report for.

If you choose **CSV** or **Excel**, choose the table or pivot table visuals from any sheet of the dashboard that you want to include in the report. You can select up to 5 visuals for each schedule.

If you choose **Excel**, one Excel workbook is generated as a final output.

8. In the **Dates** tab, choose the frequency for the report in the **Repeat** dropdown. If you're not sure, choose **Send once (Does not repeat)**.
9. For **Start date**, choose the start date and runtime that you want to send the first report on.
10. For **Timezone**, choose the time zone from the dropdown.
11. In the **Email** tab, for **E-mail subject line**, enter a custom subject line, or leave it blank to use the report title.
12. Enter the email addresses of the Quick group name of the users or groups that you want to receive the report. You can also select the **Send to all users with access** box to send the report to everyone that has access to the dashboard in your account.
13. For **Email header**, enter the header that you want the email report to show.
14. (Optional) For **E-mail body text**, leave it blank or enter a custom message to display at the beginning of the email.
15. (Optional) For PDF attachments, you can choose **Include sheet in email body** to show the first page of the PDF snapshot in the email's body.
16. Choose the method of attachment that you want the report to use. The following options are available.
 - **File attachment**– uploads an attachment of the snapshot to the email. The email size can't exceed 10 MB. This limit includes all attachments.
 - **Download link**– adds a link to the email body that users can access to download the snapshot report. When a user chooses the download link, they are prompted to sign in before the report starts to download. The link expires one year after the report is sent.

17. (Optional, recommended) To send a sample of the report before you save changes, choose **Send test report**. This option displays beside the user name of the owner of the dashboard.
18. Do one of the following:
 - (Recommended) Choose **Save** to confirm your entries.
 - To immediately send a report, choose **Save and run now**. The report is sent immediately, even if your schedule's start date is in the future.

How billing works for email reports

Authors and admins can receive any number of email reports at no extra charge.

For readers (users in the reader role), it costs one session per report, up to the monthly maximum. After receiving an email report, the reader gets a session credit to access the dashboard at no additional cost during the same month. Reader session credits don't carry over to the next billing month.

For a reader, charges for email reports and interactive sessions both accrue up to the monthly maximum charge. For readers who hit the monthly max charge, there are no further charges, and they can receive as many additional email reports as they need.

Subscribing to email reports in Amazon Quick Sight

In Enterprise edition, Amazon Quick authors can set up subscriptions to a dashboard in report form. For more information, see [Scheduling and sending Quick Sight reports by email](#). Quick readers and authors can then subscribe to a dashboard and adjust their report settings. For more information about subscribing to dashboards as a reader, see [Subscribing to Amazon Quick Sight dashboard emails and alerts](#).

Use the following procedure to change your subscription and report settings for a specific dashboard.

1. First, open a dashboard that is shared with you, or a dashboard that you own or co-own.
2. Choose the **Reports** icon at top right.
3. The **Change report preferences** screen appears. This screen shows the current report schedule, in addition to the subscription and optimization options.

For **Subscription**, choose **Subscribe** to start receiving reports, or **Unsubscribe** to stop receiving reports.

Under **Optimize**, choose the device you prefer to view the report on.

- If you usually use a mobile device or you prefer to view reports in a portrait format, choose **Viewing on a mobile device**. When you receive the report, the visuals display in a single vertical column.
- If you usually use a desktop or you prefer to view reports in a landscape format, choose **Viewing on a desktop**. When you receive the report, the visuals display in the same layout shown in your dashboard on your desktop.

4. Choose **Update** to confirm your choices, or choose **Cancel** to discard your changes.

Working with threshold alerts in Amazon Quick Sight

Applies to: Enterprise Edition

To stay informed about important changes in your data, you can create threshold alerts using KPI, Gauge, Table, and Pivot table visuals in an Amazon Quick Sight dashboard. With these alerts, you can set thresholds for your data and be notified by email when your data crosses them. You can also view and manage your alerts at anytime in a Quick Sight supported web browser.

For example, let's say that you're a customer success manager for a large organization and you want to know when the number of tickets in a support queue exceeds a certain number. Let's say too that you have a dashboard with a KPI, Gauge, Table or Pivot table visual that tracks the number of tickets in this queue. In this case, you can create an alert and be notified by email when the number exceeds the threshold you specified. That way, you can take action as soon as you're notified.

You can create multiple alerts for a single visual. If the visual is updated or deleted by the author after you create an alert, your alert settings don't change. When you create an alert, the alert takes on any filters applied to the visual at that time. If you or the author changes the filter, your existing alert doesn't change. However, if you create a new alert, your new alert takes on the new filter settings.

For example, let's say you have a dashboard with a filter control that you can use to switch the data for each visual in the dashboard from one US city to another. You have a KPI visual on the dashboard that shows average flight delays, and you're interested in delays for flights leaving from Seattle, Washington, in the US. You change the filter control to Seattle and set an alert on

the visual. This alert tracks flight delays from Seattle. Tomorrow, let's say that you want to also track flight delays from Portland, Oregon, so you change the filter control to Portland and create another alert. This new alert tracks flight delays from Portland. You now have two alerts, one on Seattle and one on Portland, working independently.

Threshold alerts are not available in the eu-central-2 Europe (Zurich) region.

For more information on KPI, Gauge, Table, or Pivot table visuals, see [Visual types in Amazon Quick Sight](#).

Note

You can't create alerts for visuals in an embedded dashboard or from the Quick mobile app. For table visuals, threshold alerts can't be created for values that are located in the Group by field well. Alerts can only be created for values that are located in the Value field well. KPI visuals that don't use a date-time field as a trend don't support alerts. An example is a KPI that shows the difference in flights between carriers X and Y instead of a KPI that shows the difference in flights between dates A and B.

Use the sections below to create and configure threshold alerts for KPI, Gauge, Table, and Pivot table visuals in Quick Sight.

Topics

- [Alert Permissions](#)
- [Creating Alerts](#)
- [Managing Threshold Alerts](#)
- [Investigating Alert Failures](#)
- [Alert Scheduling](#)
- [Using Quick connectors in threshold alerts](#)

Alert Permissions

If you're an administrator, you can control who in your organization can set threshold alerts in Quick Sight by creating a custom permissions policy. To set custom permissions in Quick, choose your user name at the upper-right corner of any Quick page, choose **Manage Quick**, and then choose **Custom permissions**.

Creating Alerts

Use the following procedure to create threshold alerts for KPI or Gauge visuals in a dashboard.

To create an alert

1. Open Quick and navigate to the dashboard that you want.

For more information about viewing dashboards as a dashboard subscriber in Quick, see [Interacting with Amazon Quick Sight dashboards](#).

2. In the dashboard, select the visual that you want to create an alert for, open the menu at the upper-right, and choose **Create alert**.
3. On the menu at upper-right on the visual, choose the **Create alert** icon.

Alternatively, you can choose the alert icon in the blue toolbar at upper right. Then, in the **Create alert** page that opens, select the KPI, Gauge, Table or Pivot table visual that you want to create an alert for, and then choose **Next**.

You can also create alerts on table or pivot table visuals by selecting a cell and choosing **Create alert**. You can only create alerts for single cells. Alerts can't be created for entire columns or for values that use a custom aggregation. For more information about custom aggregations, see [Aggregate functions](#).

4. On the **Create alert** page that opens at right, do the following:

- a. For **Name**, enter a name for the alert.

By default, the visual name is used for the alert name. You can change it if you want.

- b. For **Value to track**, choose a value that you want to set the threshold for. The information presented will vary based on the visual type you're creating an alert for.

The values that are available for this option are based on the values the dashboard author sets in the visual. For example, let's say you have a KPI visual that shows a percent difference between two dates. Given that, you see two alert value options: percent difference and actual.

If there is only one value in the visual, you can't change this option. It is the current value and it is displayed here so that you can use it as a reference while you choose a threshold. For example, if you're setting an alert on average cost, this value will show you what the

current average cost is (say, \$5). With this reference value you can make more informed decisions while setting your threshold.

- c. For **Condition**, choose a condition for the threshold.

You can choose the following conditions.

- **Is above** – Sets a rule that the alert triggers if the alert value goes above the threshold you set.
- **Is below** – Sets a rule that the alert triggers if the alert value goes below the threshold that you set.
- **Is equal to** – Sets a rule that the alert triggers if the alert value is equal to the threshold you set.

- d. For **Threshold**, enter a value to prompt the alert.

- e. For **Notification preference**, choose how often you want to be notified about a breach to the threshold you set.

You can choose from the following options.

- **As frequently as possible** - This option alerts you whenever the threshold is breached. If you choose this option, you might get alerts multiple times a day.
 - **Daily at most** - This option alerts you once per day when the threshold is breached.
 - **Weekly at most** - This option alerts you once per week when the threshold is breached.
- f. (Optional) Choose **Email me when there is no data** - When you select this option, you're notified when there's no data to check your alert rule against.
- g. Choose **Save**.

A message at upper-right appears indicating that the alert has been saved. If your data crosses the threshold you set, you get a notification by email at the address that's associated with your Quick account.

Managing Threshold Alerts

You can edit your existing alerts, turn them on or off, or view the history of times when the alert was triggered. Use the following procedures to do so.

To edit an existing alert

1. Open Quick, choose **Dashboards**, and then navigate to the dashboard that you want to edit an alert for.
2. On the Dashboards page, choose **Alerts** at upper-right.
3. On the **Manage alerts** page that opens, find the alert that you want to edit, and then choose **Edit** beneath the alert name.

You can edit the alert name, condition, and threshold.

4. Choose **Save**.

To view the history of when an alert was triggered

1. Open Quick, choose **Dashboards**, and then navigate to the dashboard that you want to view alert history for.
2. On the Dashboards page, choose **Alerts** at upper-right.
3. On the **Manage alerts** page that opens, find the alert that you want to view the history for, and then choose **History** beneath the alert name.

To turn on or turn off an existing alert

1. Open Quick, choose **Dashboards**, and navigate to the dashboard that you want to turn on or turn off an alert for.
2. On the Dashboards page, choose **Alerts** at upper-right.
3. On the **Manage alerts** page that opens, find the alert that you want to turn on or off, and then select or clear the toggle by the alert name.

The alert is turned on when the toggle is blue, and turned off when the toggle is gray.

To delete an existing alert

1. Open Quick, choose **Dashboards**, and navigate to the dashboard that you want to delete an alert from.
2. On the Dashboards page, choose **Alerts** at upper-right.
3. On the **Manage alerts** page that opens, find the alert that you want to turn on or off, choose the three-dot menu next to the alert, and then choose **Delete** from the dropdown.

Investigating Alert Failures

When an alert fails, Quick sends you an email notification about the failure. Alerts can fail for many reasons, including the following:

- The dataset the alert is using was deleted.
- The owner of the alert lost permissions to the dataset or to certain rows or columns in the dataset.
- The owner of the alert lost access to the dashboard.
- There is no data for the data tracked by the alert.

When a failure occurs, Quick sends you a notification and disables the alert if the reason for the failure isn't likely to be fixed. For example, if the alert fails due to the loss of access to a dashboard, or if the dashboard was deleted. Otherwise, Quick attempts to check your data for threshold breaches again. After four failures, Quick turns off the alert and notifies you that the alert is turned off. If the alert can be checked again, Quick sends you a notification.

To investigate why an alert failed, check that you still have access to the dashboard. Also check that you have permissions to the correct dataset and to the correct rows and columns in the dataset. If you have lost access or permissions, contact the dashboard owner. If you have the necessary access and permissions, you might need to edit your alert to avoid future alert failures.

Alert Scheduling

When you create an alert, Quick checks your data for any breaches against the thresholds you set based on when your dataset is scheduled to refresh. The information presented in the alert varies based on the visual type that you're creating an alert for. For SPICE datasets, alert rules are checked after a successful refresh of your SPICE dataset. For direct query datasets, alert rules are checked at a random time between 6:00 PM and 8:00 AM in the AWS Region that holds the dataset by default.

If you're a dataset owner, you can set an alert evaluation schedule in the dataset settings. See the following procedure to learn how.

To set an alert evaluation schedule for a dataset

1. In Quick, choose **Data** in the navigation bar at left.
2. Choose the dataset that you want to schedule alert evaluations for.

3. Choose **Set alert schedule**.
4. In the **Set alert schedule** page that opens, do the following.
 - For **Time zone**, choose a time zone.
 - For **Repeats**, choose how often you want the data to be evaluated.
 - For **Starts**, enter the time that you want the alert evaluation to start.

Using Quick connectors in threshold alerts

Prerequisites

Before you begin, make sure to [create at least one connector](#).

The connector must meet these requirements:

- Uses the **Service Auth** authentication method
- Uses one of the following integrations:
 - Atlassian Jira Cloud
 - Microsoft Outlook
 - Salesforce
 - ServiceNow

Enable Quick actions on a dashboard to use connectors

To enable Quick actions on a dashboard to use connectors

1. If a dashboard exists, go to the source analysis of the dashboard. Otherwise, [create a new analysis](#).
2. Choose **Publish**.
3. Choose between **New Dashboard** or **Replace existing dashboard**.
4. Choose the **Enable Quick actions** checkbox under **Dashboard options**.
5. Choose **Publish dashboard**.

Use connectors in a threshold alert

To use connectors in a threshold alert

1. Open a dashboard with the **Enable Quick actions** publishing option turned on.
2. Hover over a visual that supports threshold alerts. The types of visuals that support alerts can be found [here](#).
3. Choose the bell icon.
4. The **Create alert** pane opens at right.
5. Choose **Add Action**.
6. A menu appears with a list of all supported connectors and actions.
7. Choose the desired action from the list.
8. An **Action** form appears in the right pane.
9. Enter all the information you need to include with the action.
10. Some fields allow the inclusion of autofill values. Choose **Autofill** to open the menu. Choose the values you need and they will be added to your entered text.
 - **Value:** Injects the current value that was used by the alert to evaluate the alert condition
 - **Alert name:** Injects the alert name
 - **Condition:** Injects the alert condition
 - **Threshold:** Injects the threshold value
 - **All:** Injects all of the above
11. Some actions support the ability to include an attachment. You can optionally attach a PDF of the current dashboard sheet with these actions by selecting the **Include this sheet as PDF** checkbox.
12. Choose **Add action** to add the action to the alert.
13. Back at the **Create alert** pane, the configured action is added to the alert at the bottom.
14. Configure all the other desired fields of the alert and choose **Save**.
15. When your configured threshold is breached, this action should get invoked. To learn more about when a threshold alert is evaluated, see [Alert Scheduling](#).

Security and customizations

Custom Permissions/Capability Customization

- **Actions** capability: You cannot see or use actions if your user or role is restricted the permission to use the Actions capability
- **Export To PDF** capability:
 - New actions on alerts: You will not see the option to attach a PDF of the sheet while adding a new action to an alert if your user or role is restricted the permission to use the **Export To PDF** capability.
 - Existing actions on alerts: If you have existing alerts with actions containing PDF attachments then those actions will be sent out without the PDF attachments when your user or role gets restricted from using the **Export To PDF** capability.

To learn more about custom permissions, see [Custom permissions in Amazon Quick](#).

Row Level Security (RLS) / Column Level Security (CLS)

- New actions on alerts: If your dashboard contains a dataset with RLS or CLS, then
 - You cannot add actions to new alerts that track the dataset with RLS or CLS
 - You can add actions to new alerts that track a different dataset without RLS or CLS, but you cannot include PDF attachments in these actions
- Existing actions on alerts: If you add RLS or CLS to a dataset after creating alerts with actions, then
 - Existing actions on alerts tracking that dataset will stop working completely
 - Existing actions on alerts tracking a different dataset on the same dashboard will be sent out without any PDF attachments

To learn more about RLS, see [Using row-level security in Amazon Quick](#).

To learn more about CLS, see [Using column-level security to restrict access to a dataset](#).

Dashboard publishing options

- Enable PDF generation for interactive sheets
 - New actions on alerts: You will not see the option to attach a PDF of the sheet while adding a new action on alert if your dashboard has the **Enable PDF generation for interactive sheets** publishing option disabled.

- Existing actions on alerts: If you have existing alerts with actions containing PDF attachments, then those actions will be sent out without the PDF attachments when the **Enable PDF generation for interactive sheets** publishing option gets disabled on your dashboard.
- Enable Quick actions
 - New actions on alerts: You will not see the option to add an action to the alert if your dashboard has the **Enable Quick actions** publishing option turned off.
 - Existing actions on alerts: Your existing actions on alerts will stop working completely when the **Enable Quick actions** publishing option gets disabled on your dashboard.

To learn more about dashboard publishing options, see [Publishing dashboards](#).

Printing a dashboard or analysis

You can print a dashboard or an analysis in Amazon Quick Sight.

Use the following procedure to print.

1. Open the dashboard or the analysis that you want to print.
2. Choose the **Print** icon at top right.
3. On the **Prepare for printing** screen, choose the paper size and orientation that you want to use.
4. Choose **Go to Preview**.
5. Do one of the following:
 - To proceed to printing, choose **Print** to open your operating system's print dialog.
 - To make changes to the paper size or orientation, choose **Configure**.
6. To exit the preview screen, choose **Exit preview**.

Exporting Amazon Quick Sight analyses or dashboards as PDFs

You can export content from a dashboard into a Portable Document Format file (PDF). Similar to a print-out, this format provides a snapshot of the current sheet as it appears on-screen at the time of download.

To export a dashboard sheet as a PDF

1. Open Quick and choose **Dashboards** on the navigation pane at left.
2. Open the dashboard that you want to export.
3. At upper right, choose **Export, Download as PDF**. The download is prepared in the background.

When the file is ready to download, a message appears saying **Your PDF is ready..**

4. Choose **Download now** to download the file. Choose **Close** to close without downloading.

If you close this dialog box without downloading the file and want to recreate the file, repeat the previous step. Also, the downloadable file is available only temporarily for five minutes. If you wait too long to download it, the file expires. If this happens, Quick Sight instead displays an error message saying that the request has expired.

5. Repeat the previous steps for each sheet that you want to export.

You can also attach PDFs to dashboard email reports. For more information, see [Scheduling and sending Quick Sight reports by email](#).

Error codes for failed PDF export jobs

When you generate PDF reports in Amazon Quick Sight, you may encounter instances where your request to generate a PDF report fails. There are many reasons why a failure might occur. Quick Sight provides error codes that can help you understand why the error occurred and provide guidance to troubleshoot the issue. The following table lists the error codes that Quick Sight returns when a PDF export job fails.

Error code	Guidance
INVALID_DATAPREP_SYNTAX	Check the syntax for your calculated fields, and try again.
POST_AGGREGATED_METRIC_AS_DIMENSION	Aggregated metrics/operands can't be used as visual's grouping dimensions. Choose a valid visual's grouping dimensions, and try again.

Error code	Guidance
SPICE_TABLE_NOT_FOUND	The dataset has been deleted or is unavailable. Import a valid dataset, and try again.
FIELD_NOT_FOUND	A field is no longer available. Update or replace the missing fields in this dataset, and try again.
FIELD_ACCESS_DENIED	You don't have access to some fields in this dataset. Request access, and try again.
PERMISSIONS_DATASET_INVALID_COLUMN_VALUE	An invalid row level permission column value was found. Check your parent dataset rules, and try again.
COLUMN_NOT_FOUND	Replace the missing columns in your filters or parameters, and try again.
INVALID_COLUMN_TYPE	Some fields' data types have been changed and can not be automatically updated. Adjust these fields in your dataset, and try again.
PERMISSIONS_DATASET_USER_DENIED	You don't have access to this dataset. Request access to this dataset, and try again.
DATA_SOURCE_TIMEOUT	Your query has timed out. Reduce the amount of data, or import the data into SPICE, and try again.
MAX_PAGE_EXCEEDED_ERROR	Your file is ready but content is not complete. PDFs have a 1,000 page limit. Choose a shorter PDF, and try again.
INSUFFICIENT_BODY_HEIGHT_ERROR	Adjust the header and footer to be less than the page height, and try again.
FIRST_PAGE_HEIGHT_TOO_SMALL_ERROR	Adjust sections to make room for your tables, and try again.

Error code	Guidance
INTERNAL_ERROR	We can't create your PDF right now. Wait a few minutes, and try again.

Organizing assets into folders for Amazon Quick

Applies to: Enterprise Edition

In Quick Enterprise edition, your team members can create personal and shared folders to add hierarchical structure to Quick asset management. Using folders, people can more easily organize, navigate through, and discover dashboards, analyses, datasets, data sources, and topics. Within a folder, you can still use your usual tools to search for assets or to add assets to your favorites list.

You can use the following types of folders with Quick:

- Personal folders to organize work for yourself.

Personal folders are visible only to the person who owns them. You can't transfer ownership of personal folders to anyone else.

- Shared folders:
 - **Shared folders** organize work and simplify sharing among multiple people. To create and manage shared folders, you need to be a Quick administrator.
 - **Shared restricted folders** are a type of shared folder in Quick that ensure that assets remain in the shared folder. Assets that are created from assets that exist within a shared restricted folder must also stay in the restricted folder. Assets that are located in restricted folders can't be moved or shared outside of the restricted folder. For example, if you create a dataset that uses a data source that's located in a shared restricted folder, the new dataset can't be moved outside of that shared restricted folder.

Assets that are located in a restricted folder can be moved within the restricted folder tree into one or more subfolders. Subfolders of restricted folders behave like restricted folders, but dependent assets can exist in different subfolders under the same root restricted folder. The root restricted folder acts as a boundary that all assets in all subfolders can exist in as long as they remain within the root folder tree. For example, a dataset that is located in one subfolder

can use a data source that is located either another subfolder in the same folder tree or in the root folder. Any supported asset type can be created in a root folder or in any of its subfolders. Users can have different roles in different subfolders. Subfolder permissions are inherited from the parent folders of that subfolder.

Restricted folder boundary enforcement scope

Restricted folders enforce the security boundary only between these assets: analyses, dashboards, datasets, data sources, and topics. Other asset types and features in Quick can access data in a restricted folder.

For example, an analysis in a restricted folder can be built only on a dataset within the same restricted folder tree, so it respects the boundary. However, a chat agent not in a restricted folder can still retrieve details from a dataset in a restricted folder.

Restricted folders can only be created with the Quick [CreateFolder](#) API operation.

- Users that are viewers on a folder and have the Author or Admin role in Quick can view all asset types that are in the folder. Users that are viewers on a folder and have the Reader role in Quick can only see dashboards and stories that are in the folder.

All shared folders are visible to people who have access to them.

Use the following topics to learn more about creating and configuring a folder or subfolder in Quick.

Topics

- [Considerations for Quick folders](#)
- [Overview of Quick folders](#)
- [Permissions for Quick shared folders](#)
- [Create and manage membership permissions for Quick shared folders](#)
- [Creating Quick scaled folders with the Quick APIs](#)

Considerations for Quick folders

Before you get started creating and modifying folders in Amazon Quick, review the following limitations that apply to Quick folders.

- You can't share folders in your AWS account with people in other AWS accounts.
- For people who have Quick reader permissions, the following limitations apply:
 - Readers can't own a personal or shared folder.
 - Readers can't create or manage folders or folder content.
 - Readers can't have the *contributor* access level.
 - In shared folders, readers can only see dashboard assets.

In addition, these limitations apply specifically to shared folders:

- The name of a shared folder (at the top level of the tree) must be unique in your AWS account.
- In a single folder, multiple assets can't have the same name. For example, in your top-level folder, you can't create two subfolders with the same name. In the same folder, you can't add two assets with the same name, even if they have different asset IDs. The path to each asset behaves like an Amazon S3 key name. It must be unique in your AWS account.
- Restricted shared folders can only be created with the Quick CLI.

See [Overview of Quick folders](#) to learn more about the different types of folder available in Amazon Quick.

Overview of Quick folders

In Quick, you can create personal and shared folders. You can also favorite your personal or shared folders for quick access by choosing the favorite (



)

icon next to it.

You can do the following with personal folders:

- Create subfolders.
- Add assets to your folder, including analyses, dashboards, datasets, and data sources. To add assets to a personal folder, you must already have access to the assets. Multiple assets can have the same name.

Shared folders (unrestricted)

Quick administrators can perform the following tasks with shared folders.

- Create or delete a shared folder and subfolders inside of it. You can move either of these around within the top-level folder.
- Add or remove owners, contributors, and viewers. When you make a person an *owner* of the folder, you give them ownership of every asset in the folder. For more information, see [Permissions for Quick shared folders](#).

The following table summarizes the actions that a Quick user can take when working with unrestricted shared folders based on their role.

Action	Owner	Contributor	Viewer
Share an asset in a folder with users that don't have access to the folder	✔ Yes	✘ No	✘ No
Modify folder permissions	✔ Yes	✘ No	✘ No
Create assets in the folder	✔ Yes	✔ Yes	✘ No
Modify assets in the folder	✔ Yes	✔ Yes	✘ No
Delete assets in the folder	✔ Yes	✔ Yes	✘ No
Add an existing asset to a folder	✔ Yes	✔ Yes	✘ No
Remove an asset from a shared folder	✔ Yes	✘ No	✘ No
View assets in the folder	✔ Yes	✔ Yes	✔ Yes

Action	Owner	Contributor	Viewer
Create downstream assets outside of the shared folder that use assets that are located in the shared folder	✔ Yes	✔ Yes	✔ Yes*
Create downstream assets in the folder that use assets that are located outside of the folder	✔ Yes	✔ Yes	✘ No
Create subfolders	✔ Yes	✔ Yes	✘ No
Delete subfolders	✔ Yes	✘ No	✘ No
Manage subfolder permissions	✔ Yes	✘ No	✘ No
Add existing assets to subfolders	✔ Yes	✘ No	✘ No
Create new assets in subfolders	✔ Yes	✔ Yes	✘ No

Action	Owner	Contributor	Viewer
Delete assets in subfolders	✔ Yes	✔ Yes	✘ No

*The user must be assigned an admin or author role to create assets.

Restricted shared folders

Restricted shared folders provide an additional security boundary that restricts the sharing of data outside of the folder. Administrators with the appropriate IAM permissions can perform the following tasks with restricted shared folders.

- Restricted folders can be created using the `CreateFolder` API operation. For more information about the `CreateFolder` API operation, see [CreateFolder](#).
- The contributor role is assigned to users that can create and edit assets within the restricted folders. Contributors can't manage the permissions of the folder or of the assets that are in the restricted folder.
- Administrators can assign folder contributor and viewer permissions to users with the `UpdateFolderPermissions` API operation. For more information about the `UpdateFolderPermissions` API operation, see [UpdateFolderPermissions](#).

The following table summarizes the actions that a Quick user can take when working with restricted shared folders based on their role.

Action	Contributor	Viewer
Share an asset in a folder with users that don't have access to the folder	✘ No	✘ No
Modify folder permissions	✘ No	✘ No
Create assets in the folder	✔ Yes	✘ No

Action	Contributor	Viewer
Modify assets in the folder	✔ Yes	✘ No
Delete assets from the folder	✔ Yes	✘ No
Add an existing asset to a folder	✘ No	✘ No
Remove an asset from a shared folder	✘ No	✘ No
View assets in the folder	✔ Yes	✔ Yes
Create downstream assets outside of the shared folder that use assets that are located in the shared folder	✘ No	✘ No
Create downstream assets in the folder that use assets that are located outside of the folder	✘ No	✘ No
Create subfolders	✔ Yes	✘ No
Delete subfolders	✘ No	✘ No
Manage subfolder permissions	✘ No	✘ No
Add existing assets to subfolders	✘ No	✘ No

Action	Contributor	Viewer
Create new assets in subfolders	✔ Yes	✘ No
Delete assets from subfolders	✔ Yes	✘ No

The owner role is not supported for restricted shared folders.

After you choose which folder type best fits your use case, see [Permissions for Quick shared folders](#) and [Create and manage membership permissions for Quick shared folders](#) to create folders and set up folder permissions.

Permissions for Quick shared folders

Shared folders have three permission levels. To set folder-level permissions for a user or group, see [Create and manage membership permissions for Quick shared folders](#).

- **Owners** - The folder *owner* owns everything (folders, analyses, dashboards, datasets, data sources, topics) inside of the folder. They can create, edit, and delete the assets in the folder, modify permissions on the folder and its assets, and delete the folder entirely. The owner role is not supported for restricted shared folders.
- **Contributors** - A *contributor* can create, edit, and delete assets in a folder just like an owner. They can't delete the folder or modify permissions on the folder or on assets where they have contributor access that they inherited from the folder.
- **Viewers** - A *viewer* can only view the assets (folders, dashboards, datasets, data sources, topics) in the folder. A viewer can't edit or share those assets.

The following rules also apply to security for shared folders:

- Quick readers' sharing status for a folder gets shared with the folder. However, a reader gets only read access to folders, and only dashboard access to visuals.
- AWS security is enforced on every object within a folder. The folder applies the same type of security to the assets of whoever the folder is shared with according to their access level (admin, author, or reader).

- The *top-level folder* is the root folder of any subfolders. When a subfolder is shared at any level, the person whom the folder was shared with sees the root folder in the top-level folders view.
- The folder permission is the permission on the current folder, combined with permissions of all the folders leading to the root folder.
- A *shared asset* inherits its permission from the folder. A shared asset is created when an asset that belongs to the folder owner is added to a shared folder.
- If you own an unrestricted shared folder, you can transfer ownership of the folder to another Quick admin.
- The owner role is not supported for restricted folders. The contributor role is assigned to authors that create and edit assets within the restricted folders. Folder contributors can't manage the permissions of the restricted folder or its assets.
- The correct IAM permissions are required to update the permissions of a restricted shared folder with the `UpdateFolderPermissions` API.

To create and manage permissions of a shared folder, see [Create and manage membership permissions for Quick shared folders](#).

Create and manage membership permissions for Quick shared folders

Shared folders (unrestricted)

To create a shared folder and to share the folder with one or more groups in the Quick console, you must be an Amazon QuickSight administrator. You can also create a shared folder with the `CreateFolder` API operation. Use the following procedure to share or modify the membership permissions of a shared folder.

1. From the left navigation, choose **Folders** then **Shared Folders**. Find the folder that you want to share or manage permissions for.
2. To open the actions menu for that folder's row, choose the ellipsis (three dots).
3. Choose **Share**.
4. In the **Share folder** modal, add the groups and users with whom you want to share the contents of the folder.
5. For each user and group that you add, choose a permission level from the **Permissions** menu in that row.
6. To update the permission type for an existing user, choose **Manage folder access**.

7. When you're done setting user and group permissions for the folder, choose **Share**. Users are not notified that they now have access to the folder.

Restricted shared folders

Restricted shared folders can only be created with the CreateFolder API operation. The following example creates a restricted shared folder.

```
aws quicksight create-folder \  
--aws-account-id AWSACCOUNTID \  
--region us-east-1 \  
--folder-id example-folder-name \  
--folder-type RESTRICTED \  
--name "Example Folder" \
```

After you create a restricted shared folder, assign folder contributor and viewer permissions with a UpdateFolderPermissions API call. The following example updates the permissions of a restricted shared folder to grant contributor permissions to a user.

```
aws quicksight update-folder-permissions \  
--aws-account-id AWSACCOUNTID \  
--region us-east-1 \  
--folder-id example-folder-name \  
--grant-permissions Principal=arn:aws:quicksight::us-east-1::AWSACCOUNTID:user/default/:username,Actions=quicksight:CreateFolder  
,quicksight:DescribeFolder, \  
quicksight:CreateFolderMembership,quicksight>DeleteFolderMembership,qu  
icksight:DescribeFolderPermissions \
```

The permissions that you pass to the user depend on the type of folder role that you want to grant them. Use the following lists to determine which permissions are needed for the user that you want to grant folder access to.

Folder owner

- quicksight:CreateFolder
- quicksight:DescribeFolder
- quicksight:UpdateFolder
- quicksight>DeleteFolder

- quicksight:CreateFolderMembership
- quicksight>DeleteFolderMembership
- quicksight:DescribeFolderPermissions
- quicksight:UpdateFolderPermissions

Folder contributor

- quicksight:CreateFolder
- quicksight:DescribeFolder
- quicksight:CreateFolderMembership
- quicksight>DeleteFolderMembership
- quicksight:DescribeFolderPermissions

Folder viewer

- quicksight:DescribeFolder

After you create a shared folder, you can begin using the folder in Quick.

You can also use the Quick APIs to create special scaled folders that can be shared with up to 3000 namespaces. To learn more about creating a scaled folder, see [Creating Quick scaled folders with the Quick APIs](#).

Creating Quick scaled folders with the Quick APIs

You can use the Amazon Quick APIs to create special scaled folders that can be shared with up to 3000 namespaces. Each namespace that is added to a folder can contain up to 100 principals. A *principal* is a user or a group of users. After you create a scaled folder and add the principals that you want, any QuickSight asset can be added to the folder. It can then be shared to every principal in the namespaces that the folder principals are assigned to. This streamlines the process to share Quick assets with thousands of users.

Scaled folders can only be created with the Quick APIs. When you create a scaled folder, you can share the folder with up to 100 principals that are in the same namespace. You can add principals that belong to a different namespace with an UpdateFolderPermissions API call. After the

folder is created, you can add and remove assets from the folder with the Quick APIs or the Quick console.

Each Amazon Quick account holds up to 100 scaled folders. You can add up to 100 assets to a scaled folder. If you want to share a scaled folder with more than 3000 namespaces, contact [AWS support](#).

Examples

The following examples show how to create a scaled folder with the Quick APIs.

Prerequisites

Before you begin, verify that you have an AWS Identity and Access Management role that grants the API user access to call the Quick API operations. The following example shows an IAM policy that you can add to an existing IAM role to create, delete, or modify a scaled folder. With the sample policy, users can add dashboards, analyses, and datasets to a scaled folder.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "quicksight:CreateFolder",
        "quicksight:CreateFolderMembership",
        "quicksight>DeleteFolderMembership",
        "quicksight>DeleteFolder",
        "quicksight:DescribeFolderPermissions",
        "quicksight:DescribeFolderResolvedPermissions",
        "quicksight:UpdateFolderPermissions",
        "quicksight:UpdateDashboardPermissions",
        "quicksight:UpdateAnalysisPermissions",
        "quicksight:UpdateDataSetPermissions"
      ],
      "Resource": "*"
    }
  ]
}
```


The following example creates a scaled folder.

```
aws quicksight create-folder \  
--aws-account-id "AWSACCOUNTID" \  
--region "us-east-1" \  
--name "eastcoast-users" \  
--sharing-model "NAMESPACE" \  
--folder-id "eastcoast-users"
```

After you create a scaled folder, share the folder with a principal in your account. You can only grant or revoke permissions to users and groups that are within the same namespace in each API call. The following example shares a scaled folder with a user in the same account that the folder exists in.

```
aws quicksight update-folder-permissions \  
--aws-account-id "AWSACCOUNTID" \  
--region "us-east-1" \  
--folder-id "eastcoast-users" \  
--grant-permissions \  
  '[  
    {"Actions":  
      ["quicksight:DescribeFolder",  
       "quicksight:UpdateFolder",  
       "quicksight>DeleteFolder",  
       "quicksight:DescribeFolderPermissions",  
       "quicksight:UpdateFolderPermissions",  
       "quicksight:CreateFolderMembership",  
       "quicksight>DeleteFolderMembership",  
       "quicksight:CreateFolder"  
      ],  
      "Principal": "arn:aws:quicksight:us-east-1:AWSACCOUNTID:user/default/my-user"  
    }  
  ]'
```

After you share the folder with a new principal, validate the new folder permissions with a `describe-folder-permissions` API call.

```
aws quicksight describe-folder-permissions \  
--aws-account-id "AWSACCOUNTID" \  
--region "us-east-1" \  
--folder-id "eastcoast-users" \  

```

```
--namespace "default"
```

After you validate the new folder permissions, create a subfolder within the scaled folder. The subfolder inherits the permissions of the scaled folder that it's created in.

```
aws quicksight create-folder \  
--aws-account-id "AWSACCOUNTID" \  
--region "us-east-1" \  
--name "new-york-users" \  
--sharing-model "NAMESPACE" \  
--folder-id "new-york-users" \  
--parent-folder-arn "arn:aws:quicksight:us-east-1:AWSACCOUNTID:folder/eastcoast-users"
```

The following example validates the inherited permissions of the new subfolder.

```
aws quicksight describe-folder-resolved-permissions \  
--aws-account-id "AWSACCOUNTID" \  
--region "us-east-1" \  
--folder-id "new-york-users" \  
--namespace "default"
```

After you validate the permissions of the subfolder, add the Quick asset that you want to share to the folder. After you add the asset to the subfolder, the asset is shared with every principal that the subfolder is shared with. The following example adds a dashboard to a subfolder.

```
aws quicksight create-folder-membership \  
--aws-account-id "AWSACCOUNTID" \  
--folder-id "new-york-users" \  
--member-id "my-dashboard" \  
--member-type "DASHBOARD" \  
--region "us-east-1"
```

Exploring interactive dashboards in Amazon Quick Sight

Intended audience: Amazon Quick Dashboard subscribers or viewership

In Amazon Quick Sight, a *data dashboard* is a collection of charts, graphs, and insights. It's like a newspaper that's all about the data that you're interested in, except it has digital pages. Instead of reading it, you interact with it.

Dashboards come in a wide variety of designs, depending on what you do and the analytics that you need to do it well. Using Quick Sight, you can interact with your data on a webpage or your mobile device. If you also subscribe by mail, you can see a static preview of it.

The story told by your data reflects the expertise of the analysts and data scientists who built the dashboards. They refine the data, add calculations, find angles on the story, and decide how to present it. The publisher designs the dashboard and fills it with interactive data visualizations and controls that adjust your view. Publishers can customize the level of interactivity that you have, including filter and search options. You can interact with the active items on the screen to filter, sort, drill down, or jump to another tool.

When you view a dashboard, you're seeing the most recently received data. As you interact with the items on the screen, any changes you make change your view of the dashboard, and no one else's. Thus, your device's privacy is assured, although the publisher can tell what you looked at. After you close the dashboard, your explorations aren't preserved and neither is the data. As always, while you're an Quick reader, your monthly subscription is provided by the publishers of the dashboards at no additional cost to you.

If you're also a dashboard publisher—we call them authors, because they write reports—you can also save a copy of the dashboard for further analysis. If you find a new feature of the data that you want to publish, work with the original authors to update it. That way, everyone can see the same version of the story. However, you can also use your copy to learn how their design works or to inspire your work on something entirely new. Then, when you're finished, you can publish your analysis as a new dashboard.

To learn to set up dashboards, see [Sharing and subscribing to data in Amazon Quick Sight with dashboards and reports](#).

Topics

- [Interacting with Amazon Quick Sight dashboards](#)
- [Interacting with pixel perfect reports in Amazon Quick Sight](#)
- [Subscribing to Amazon Quick Sight dashboard emails and alerts](#)
- [Creating a reader generated report in Amazon Quick Sight](#)

- [Bookmarking views of a Amazon Quick Sight dashboard](#)

Interacting with Amazon Quick Sight dashboards

To access a dashboard that you've been invited to share, follow the instructions in the invitation email. You can also access a dashboard if it's embedded into an application or website that you already have access to.

To fit the dashboard to your screen, open the **View** menu at upper right and select **Fit to window**.

Depending on how the dashboard is configured, you can find all or some of the following elements:

- The menu bar – This displays the name of the dashboard. Also, the menu bar shows what you can do with the dashboard, including **Undo**, **Redo**, and **Reset**, on the left. As you interact with the dashboard, you can use these as tools to help you explore, knowing that you can change your view without losing anything. On the right, you can find options to **Print** the dashboard, work with **Data**, choose a different AWS **Region**, and open your **User Profile**. The user profile menu has options so you can choose the language that Amazon Quick Sight displays. It also has links to the Quick **Community** and the online documentation (**Help**).
- The dashboard sheets – If your dashboard has multiple sheets, these display as tabs across the top of the dashboard.
- The **Filter** menu – This option displays to the left of the dashboard, if the dashboard publisher allows filtering.
- The **Controls** palette – If your dashboard includes controls, you can use them to choose the options (parameters) that you want to apply to your dashboard. Sometimes a control value is selected for you, and sometimes it's set to **ALL**.
- The dashboard title – If your dashboard has a title, it is usually a larger heading. It might have some status information or instructions below it.
- The dashboard widgets – The items on the screen can include charts, graphs, insights, narratives, or images. To see them all, you might need to scroll vertically or horizontally.

Using filters on Amazon Quick Sight dashboard data

You can use filters to refine the data displayed in a visual. Filters are applied to the data before any aggregate functions. If you have multiple filters, all top-level filters apply together using AND. If the filters are grouped inside a top-level filter, the filters in the group apply using OR.

Amazon Quick Sight applies all of the enabled filters to the field. For example, suppose that there is one filter of `state = WA` and another filter of `sales >= 500`. In this case, the dataset contains only records that meet both of those criteria. If you disable one of these, only one filter applies. Take care that multiple filters applied to the same field aren't mutually exclusive.

Viewing filters

To see the existing filters, choose **Filter** on the element settings menu, then choose to view filters. The filters display in the **Applied filters** panel in order of creation, with the oldest filter on top.

Understanding filter icons in an Amazon Quick Sight dashboard

Filters in the **Applied filters** panel display icons to indicate how they are scoped and whether they are enabled.

A filter that isn't enabled is grayed out, and you can't select its check box.

One of several scope icons displays to the right of the filter name to indicate the scope set on that filter. The scope icon resembled four boxes in a square. If all boxes are filled, the filter applies to all visuals on the analysis sheet. If only one box is filled, the filter applies to the selected visual only. If some boxes are filled, the filter applies to some of the visuals on the sheet, including the one currently selected.

The scope icons match the ones that display on the filter menu when you are choosing the scope for the filter.

Viewing filter details in an Amazon Quick Sight dashboard

To see filter details, choose **Filter** at left. The filter view retains your last selection. So when you open **Filter**, you see either the **Applied filters** or the **Edit filter** view.

In the **Applied filters** view, you can choose any filter to view its details. The filters in this list can change depending on the scope of the filter, and which visual you currently have selected.

You can close the **Edit filter** view by choosing the selector on the right. Doing this resets the **Filter** view.

Filtering data during your session in Amazon Quick Sight

While your dashboard session is active, you can filter data in three ways:

1. If your dashboard has controls at the top of the screen, you can use them to filter data by choosing from a preset list of values.
2. You can use the filter icon on each widget's settings menu.
3. You can create your own filters by using the filter panel on the left side of the page. The filter icon looks like the following.

To create a filter, choose the **Filter** icon at left.

The first step is to choose which dashboard element you want to filter.

Click on the item you choose, so that a highlight appears around the selected item. Also, if any filters are already there, they display in a list. If there aren't any filters, you can add one by using the plus sign (+) near **Filters**.

Filtering options vary depending on the data type of the field you want to filter, and on the options you choose inside the filter. The following screenshot shows some of the options available for a time-range date filter.

For each filter, you can choose whether to apply it to one, some, or all dashboard elements. You can also enable or disable filters by using the check box next to the name of the filter. To delete a filter, edit it and scroll to the bottom to see the options. Remember that your filters aren't saved from one session to the next.

For more detailed information on creating filters, see [Filtering data in Amazon Quick Sight](#).

Using the elements on the Amazon Quick Sight dashboard

Each widget has a settings menu that appears when you select that widget. This menu provides options to zoom in or out, filter the data, export the data, and more. The options vary depending on what type of widget the element is.

When you choose a data point, several actions are available. You can click or tap on a data point, for example on a bar in a bar chart, on a point where the line bends on a line chart, and so on. The available options vary based on what type of item it is.

These actions are as follows:

- Focus on or exclude.

You can focus on or exclude specific data in a field, for example regions, metrics, or dates.

- Drill up or drill down.

If your dashboard contains data on which you can drill down or up, you can drill up to a higher level or drill down to explore deeper details.

- Custom URL actions.

If your dashboard contains custom actions, you can activate them by choosing a data point or by right-clicking it. For example, you might be able to email someone directly from the dashboard. Or you might open another sheet, website, or application, and send it the value you chose from this one.

- Change chart colors or specific field colors.

You can change all the chart colors to a specific color. Or you can choose a specific field value to change its color of the element it's part of.

Sorting dashboard data in Amazon Quick Sight

You can sort data in three ways:

1. You can hover over the label for the field you want to sort by, and choose the sort icon.
2. You can choose the filter icon at the upper right of one of the dashboard elements.
3. You can click or tap on the field and choose **Sort** from the context menu.

Sorting for pivot tables is different; you specify the sort order by using the column sort icon on the pivot table.

Exporting and printing interactive Amazon Quick Sight dashboard reports

You can export or print a PDF version of an interactive dashboard. You can also export some visuals in a dashboard to a CSV. Exporting an entire dashboard to a CSV is not currently supported for interactive dashboards.

Exporting data from a dashboard to a PDF

To export an interactive dashboard report as a PDF

1. From the dashboard report that you want to export, choose the **Export** icon at the top right.

2. Choose **Generate PDF**.
3. When you choose **Generate PDF**, Quick Sight will begin preparing the dashboard report for download. Choose **View downloads** in the blue pop-up to open the **Downloads** pane on the right.
4. There are two ways to download your report:
 - Choose **DOWNLOAD NOW** in the green pop-up.
 - Choose the **Export** icon at the top right, and then choose **View downloads** to view and download every report that is ready to download.

To print an interactive dashboard report

1. From the report that you want to print, choose the **Export** icon at the top right, and then choose **Print**.
2. In the **Prepare for printing** pop-up that appears, choose the paper size and orientation that you want. You can optionally choose to include the background color by selecting **Print background color**.
3. Choose **GO TO PREVIEW**.
4. In the preview window that appears, choose **PRINT**.

Exporting data from a dashboard to a CSV

Note

Export files can directly return information from the dataset import. This makes the files vulnerable to CSV injection if the imported data contains formulas or commands. For this reason, export files can prompt security warnings. To avoid malicious activity, turn off links and macros when reading exported files.

To export data from an analysis or dashboard to a comma-separated value (CSV) file, use the settings menu at the upper right of a widget. Exports only include data that currently displays in the item that you choose.

In tables and pivot tables, you can export data to a comma-separated value (CSV) file or Microsoft Excel file. You can choose to export only visible fields or all fields.

To export only visible fields to a CSV or Excel file, choose the menu at upper-right of the visual. Choose either **Export to CSV** or **Export to Excel**, and then choose **Export visible fields to CSV** or **Export visible fields to Excel**.

To export all fields to a CSV or Excel file, choose the menu at upper-right of the visual. Choose either **Export to CSV** or **Export to Excel**, and then choose **Export all fields to CSV** or **Export all fields to Excel**.

Generate an executive summary of an Amazon Quick Sight dashboard

Dashboard readers can generate executive summaries that provide a summary of all insights that Quick Sight has generated for the dashboard. Executive summaries make it easier for readers to find key insights and information about a dashboard at a glance.

When readers are viewing a dashboard that uses executive summaries, the **Executive summary** option is available in the **Build** dropdown list that is located in the top right of the Dashboard's page. Use the procedure below to generate an executive summary. If a dashboard doesn't use executive summaries, the **Executive summary** option does not appear in the **Build** dropdown list.

To generate an executive summary

1. In the dashboard that you want to work in, choose **Build**, and then choose **Executive summary**.
2. Choose **Summarize**. The executive summary is generated and the appears on the left.

Executive summaries use the data of the current dashboard sheet and visual settings. If the dashboard or visual settings are updated, a warning appears at the top of an executive summary. To refresh the executive summary of an updated dashboard, generate a new executive summary.

After an executive summary is generated, Amazon Quick readers can copy the summary to their clipboard in order to share with others, or include in a Quick Sight story. For more information about Quick Sight stories, see [Working with data stories in Amazon Quick Sight](#).

Customizing tables and pivot tables in Amazon Quick Sight

Reader customization for tables and pivot tables is enabled by default. You can change the visual to fit your analysis needs without requesting updates from the dashboard author. Your customizations are private – other readers of the same dashboard don't see your changes unless you share them.

For dashboard authors

To disable reader customization, choose **Format Visual**, choose **Interactions**, and then turn off **Reader Customization**. Republish the dashboard for the change to take effect.

You can customize tables and pivot tables in the following ways:

- **Sort columns** – Organize data in ascending or descending order.
- **Reorder columns** – Rearrange columns to reflect the order that matters most to you.
- **Hide and show columns** – Focus on relevant data by hiding columns you don't need, and show them again when you do.
- **Freeze columns** – Keep important columns visible while scrolling horizontally through large datasets.
- **Add and remove fields** – Include additional fields from the dataset or remove fields you don't need.
- **Change aggregations** – Modify how a measure is aggregated (for example, change from *Sum* to *Average*).
- **Modify formatting** – Adjust field formatting directly in the dashboard view.

Note

Reader customization is supported for tables and pivot tables only. Other visual types don't support reader-level customization at this time.

Sorting columns

To sort data in a table or pivot table, choose the column header that you want to sort by. Choose it again to toggle between ascending and descending order.

Reordering columns

To rearrange columns, choose the column header menu and then choose **Move left** or **Move right**.

Hiding and showing columns

To hide a column, choose the column header menu and then choose **Hide**.

To show hidden columns, choose any column header menu and then choose **Show all hidden fields**.

Freezing columns

To freeze a column so that it stays in place while you scroll horizontally, choose the column header menu and then choose **Freeze column**.

This is useful for keeping key identifiers, such as region names or account numbers, visible while you review a wide table.

Adding and removing fields

If the author has made additional fields available for customization, you can add or remove them from the visual.

To add or remove fields

1. On the table or pivot table, choose **Customize**.
2. In the field list, select the fields you want to add (for example, *City*, *Profit*, or *Quantity*).
3. To remove a field, clear its selection in the field list.

The available fields are determined by the author. By default, you can add back, remove, hide, show, reorder, and change aggregations for the fields that are already in the visual. Authors can extend this list to include additional fields from the underlying dataset.

Changing aggregations

After you add or select a measure field, you can change its aggregation type. For example, you can change *Order Date* to aggregate by **Quarter**, or change *Quantity* from **Sum** to **Average**.

To change an aggregation, choose the field in the customization panel and then select a different aggregation type.

Resetting to the default view

To discard all of your customizations and return to the author's original configuration, choose any column header menu and then choose **Reset visual**.

Saving your customizations

Your customizations are saved automatically. When you return to the dashboard, your personalized view is preserved – you don't need to reapply settings each time you open the dashboard.

Sharing customized views

You can share your customized view with other readers in the following ways:

- **Share this view** – Generate a link that preserves your current filters, column selections, and ordering. Other users who open the link see the same view. This is useful for ad-hoc collaboration.
- **Bookmarks** – Save your customizations as a bookmark for recurring use. Bookmarks capture visual customizations and applied filters, so you can return to your preferred view at any time. Bookmarks can be private or shared across teams.

Exporting customized views

You can schedule and export your customized table or pivot table in the following formats:

- PDF
- CSV
- Excel

This is useful for sharing data with stakeholders who don't have Amazon Quick Sight access or for offline analysis.

Embedding behavior

When tables and pivot tables are embedded in an application, customization availability and persistence depend on the embedding method.

- **Visual embedding (registered or anonymous users)** – You can customize the visual. Customizations are not persisted – the original dashboard is displayed when the page reloads.
- **Dashboard embedding for registered users** – You can customize the visual. If state persistence is enabled through embedding options, your customized view is preserved on reload. If state persistence is not enabled, the original dashboard is displayed.
- **Dashboard embedding for anonymous users** – You can customize the visual. Customizations are not persisted – the original dashboard is displayed when the page reloads.

The `createSharedView` SDK function supports generating a shared view from a customized embedded dashboard.

Limitations

- Reader customization is supported for tables and pivot tables only. Other visual types, such as bar charts, line charts, and KPIs, don't support reader-level customization.
- The fields available for readers to add or remove are controlled by the dashboard author. If you need access to a field that isn't available, contact the dashboard author.

Interacting with pixel perfect reports in Amazon Quick Sight

To access a pixel perfect report that you've been invited to share, follow the instructions in the invitation email. You can also access a pixel perfect report if it's embedded into an application or website that you already have access to.

To fit the pixel perfect report to your screen, open the **View** menu at upper right and select **Fit to window**. You can also zoom in and out using the plus (+) and minus (-) icons on the top left corner of the report.

Exporting and printing Amazon Quick Sight reports

Pixel perfect reports are designed to be viewed from a specific point of time. These reports, or snapshots, can be printed or downloaded as a PDF or CSV.

To export a pixel perfect report report as a PDF

1. From the pixel perfect report that you want to export, choose the **Export** icon at the top right.
2. Choose **Generate PDF**.
3. When you choose **Generate PDF**, Quick Sight will begin preparing the pixel perfect report for download. When the report is ready, a green pop up will appear that says **Your PDF is ready**.
4. There are two ways to download your report:
 - Choose **DOWNLOAD NOW** in the green pop-up.
 - Choose the **Export** icon at the top right, and then choose **View downloads** to view and download every report that is ready to download.

To export a pixel perfect report as a CSV

1. From the report that you want to export, choose the **Scheduling** icon at the top right, and then choose **Recent snapshots**.
2. In the **Recent snapshots** menu that appears on the right, snapshots are sorted from most recently generated to the oldest. Snapshots are stored for up to 1 year. Find the report that you want to download and choose the download icon to the right of the report.
3. In the report pop-up that appears, choose the download icon next to the version of the report that you want to download. You can choose to download the report as a CSV, or you can download the report as a PDF.

To print a pixel perfect report

1. From the report that you want to print, choose the **Export** icon at the top right, and then choose **Print**.
2. When you choose **Print**, your browser's printer pop-up appears. From here, you can print the PDF the same way you would print anything else on your browser.

Subscribing to Amazon Quick Sight dashboard emails and alerts

Using Amazon Quick Sight, you can subscribe to updates for certain events, such as dashboard updates and anomaly alerts.

Topics

- [Sign up for dashboard emails](#)
- [Sign up for anomaly alerts](#)

Sign up for dashboard emails

You can sign up to get a dashboard in report form, and receive it in an email. You can also configure your report settings.

To change subscription and report settings for a dashboard

1. Open a dashboard that is shared with you.
2. Choose the **Schedules** icon at upper right, and then choose **Schedules** in the dropdown.

3. The **Schedules** pane appears on the right. This pane shows all of the different scheduled reports that you are or can be subscribed to. Navigate to the report that you want and toggle the switch to subscribe or unsubscribe from the report.

Sign up for anomaly alerts

On a dashboard that has a narrative insight that's configured for anomaly detection, you can sign up to get alerts for anomalies and contribution analysis. You receive anomaly alerts when anomalies are updated. The alerts email displays the total number of anomalies, and provides detail on the top five, according to your personal alert configuration. You receive key driver contribution analysis when it's updated, provided that contribution analysis is configured to run with anomaly detection.

To set up anomaly alerts

1. Open a dashboard that is shared with you.
2. You can configure alerts from one of two screens. Choose one of the following, then go to the next step:
 - In the dashboard, locate the anomaly widget that you're interested in. Select it so that it has a highlighted box around it.
 - If you're in the dashboard and have the **Explore Anomalies** page open, you can configure the alert without returning to the dashboard view.
3. At upper right, choose **Configure alert**. The **Alert** configuration screen appears.
4. For **Severity**, choose the lowest level of significance that you want to see.

For **Direction**, choose to get alerts about anomalies that are **Higher than expected** or **Lower than expected**. You can also choose **[ALL]** to receive alerts about all anomalies.

5. Choose **OK** to confirm your choices.
6. To stop receiving to an anomaly alert, locate the anomaly widget in the dashboard and use the bell icon to unsubscribe. You can also use the **To manage this alert** link at the bottom of an alert email.

Creating a reader generated report in Amazon Quick Sight

If a Amazon Quick author has set up a prompted report for a Quick Sight pixel perfect report, Quick Sight dashboard viewers can use the prompt to schedule their own reports for themselves. For more information about prompts for pixel perfect reports, see [Setting up prompts for paginated reports](#).

Use the following sections to learn how to create and modify a reader generated report.

Topics

- [Creating a reader generated report](#)
- [Loading a saved view of a Quick Sight reader generated report](#)
- [Updating the view of a scheduled reader generated report](#)
- [Updating a reader generated report schedule](#)

Creating a reader generated report

Use the following procedure to create a reader generated report.

To create a reader generated report

1. Open the [Quick console](#).
2. Open the dashboard that you want to create a report for.
3. Choose the **Scheduling** at the top of the dashboard page.
4. The scheduling pane opens. To add a new report schedule, choose **Add**. If you do not see the **Add** button, the dashboard does not contain a pixel perfect sheet, or your Quick account does not have the Pixel perfect reports add on. For more information about the paginated reports add on, see [Getting started](#).
5. For **Schedule name**, enter a name for the new schedule. The schedule name can be up to 100 characters long.
6. For **Description**, choose the view option that you want the report to use. You can choose from the following views:
 - **Custom view** – The current view of the dashboard.
 - **Original view** – The author published view of the dashboard.
7. For **Content**, choose the pixel perfect report sheet that you want to generate a PDF report for.

8. For **Dates**, choose the frequency at which you want to receive the report. Scheduling options that are available for an email report include the following:
 - **Once (Does not repeat)** – Sends the report only once at the date and time that you choose.
 - **Daily** – Repeats daily at the time that you choose.
 - **Weekly** – Repeats each week on the same day or days at the time that you choose. You can also use this option to send reports in weekly intervals, such as every other week or every three weeks.
 - **Monthly** – Repeats each month on the same day of the month at the time that you choose. You can also use this option to send reports on specific days of the month, such as the second Wednesday or the last Friday of each month.
 - **Yearly** – Repeats each year on the same day of the month or months selected at the time that you choose. You can also use this option to send reports on specific days or sets of days in selected months. For example, you can configure a report to be sent on the first Monday of January, March, and September, or on July 14th, or on the second day of February, April, and June each year.
 - **Custom** – Configure your own scheduled report that best fits your business needs.

The scheduled report is sent within 1 hour from the specified time. Delays may occur during peak hours.

9. In the **Email** tab, for **E-mail subject line**, enter a custom subject line, or leave it blank to use the report title.
10. Enter the email addresses of the Quick group name of the users or groups that you want to receive the report.
11. For **Email header**, enter the header that you want the email report to show.
12. (Optional) For **E-mail body text**, leave it blank or enter a custom message to display at the beginning of the email.
13. (Optional, recommended) To send a sample of the report before you save changes, choose **Send test report**.
14. Do one of the following:
 - (Recommended) Choose **Save** to confirm your entries.
 - To immediately send a report, choose **Save and run now**. The report is sent immediately, even if your schedule's start date is in the future.

After you save a report schedule, the schedule appears in the **Schedules** pane. Reader generated reports are only available to the user that created them and can't be shared.

Loading a saved view of a Quick Sight reader generated report

Amazon Quick readers can use the **Schedules** pane to load a saved view of any scheduled pixel perfect report they have created or received. Use the following procedure to load a saved review of a scheduled report.

To load a saved view of a scheduled report

1. Open the [Quick console](#).
2. Open the dashboard that contains the report that you want to change.
3. Choose the **Scheduling** at the top of the dashboard page.
4. The scheduling pane opens. Locate the schedule that you want to change and choose the ellipsis (three dots) icon next to the report to open the schedule menu, and then choose **Details**.
5. Choose **Load saved view**. The saved view of the dashboard that was used for the selected schedule is rendered. All filter values that were active when the dashboard snapshot was taken are applied to the dashboard. When a saved view of a dashboard is loaded, the reader's current view of the dashboard is lost.

Updating the view of a scheduled reader generated report

After a Amazon Quick reader has created a report in Quick Sight, they can use the **Schedules** pane to update the dashboard view that is used in the scheduled report. Use the following procedure to update the dashboard view of a scheduled report.

To change the dashboard view of a scheduled report

1. Open the [Quick console](#).
2. Open the dashboard that contains the report that you want to change.
3. Choose the **Scheduling** at the top of the dashboard page.
4. The scheduling pane opens. Locate the schedule that you want to change and choose the ellipsis (three dots) icon next to the report to open the schedule menu, and then choose **Details**.

5. Choose **Load saved view**. The saved view of the dashboard that was used for the selected schedule is rendered. All filter values that were active when the dashboard snapshot was taken are applied to the dashboard. When a saved view of a dashboard is loaded, the reader's current view of the dashboard is lost.
6. Update the dashboard filters that you want to change.
7. Choose the **Scheduling** at the top of the dashboard page.
8. The scheduling pane opens. Locate the schedule that you want to change and choose the ellipsis (three dots) icon next to the report to open the schedule menu, and then choose **Edit**.
9. Navigate to the **Dashboard view** section, and then choose **Custom view**. The new filter values that you updated are applied to the dashboard report.
10. Choose **Save** to update the schedule.

Updating a reader generated report schedule

After they create a reader generated report, Amazon Quick readers can use the **Schedules** pane to make a report schedule active or inactive. Use the following procedure to update active status of a reader generated report schedule.

1. Open the [Quick console](#).
2. Open the dashboard that contains the report that you want to change.
3. Choose the **Scheduling** at the top of the dashboard page to open the **Schedules** pane.
4. Choose **Schedules**.
5. Navigate to the **My schedules** section and find the schedule that you want to update.
6. Use the toggle to set the report schedule to **Active** or **Inactive**.
7. When you are finished making changes to the report schedule, close the **Schedules** pane.

Bookmarking views of a Amazon Quick Sight dashboard

When you load a dashboard as an Amazon Quick reader or author, you can create bookmarks to capture specific views of your interests. For example, you can create a bookmark for a dashboard with a specific filter setting that differs from the original published dashboard. By doing this, you can quickly return to the data that's relevant to you.

After you create a bookmark, you can set it as the default view of the dashboard that you see when you open the dashboard in a new session. This doesn't affect anyone else's view of the dashboard.

You can create up to 200 bookmarks for a dashboard and share them by a URL link with other subscribers of that dashboard.

Dashboard bookmarks are available on the Quick console.

Dashboard bookmarks for pixel perfect reports are currently not supported. For more information on pixel perfect reports, see [Working with pixel perfect reports in Amazon Quick Sight](#).

Use the following topics to learn how to use bookmarks.

Topics

- [Creating bookmarks in Amazon Quick Sight](#)
- [Updating bookmarks in Amazon Quick Sight](#)
- [Renaming bookmarks in Amazon Quick Sight](#)
- [Making a bookmark the default view in Amazon Quick Sight](#)
- [Sharing bookmarks in Amazon Quick Sight](#)
- [Deleting bookmarks in Amazon Quick Sight](#)

Creating bookmarks in Amazon Quick Sight

Use the following procedure to create a bookmark for a dashboard.

To create a bookmark for a dashboard

1. Open the published dashboard that you want to view and make changes to the filters or parameters, or select the sheet that you want. For example, you can filter to the Region that interests you, or you can select a specific date range using a sheet control on the dashboard.
2. Choose the bookmark icon at upper right, and then choose **Add bookmark**.
3. In the **Add a bookmark** pane that opens, enter a name for the bookmark, and then choose **Save**.

The bookmark is saved, and the dashboard name updates with the bookmark name (at top left).

You can return to the original dashboard view that the author published at any time by selecting **Original dashboard** in the **Bookmarks** pane at right.

Updating bookmarks in Amazon Quick Sight

At any time, you can change a bookmark dashboard view and update the bookmark to always reflect those changes.

To update a bookmark

1. Open the published dashboard and make needed changes to the filters or parameters, or select a sheet.
2. Choose the bookmark icon at upper right.
3. In the **Bookmarks** pane that opens, choose the context menu (the three vertical dots) for the bookmark that you want to update, and then choose **Update**.

A message appears, confirming the update.

Renaming bookmarks in Amazon Quick Sight

Use the following procedure to rename a bookmark.

To rename a bookmark

1. In a published dashboard, choose the bookmark icon at upper right to open the **Bookmarks** pane.
2. In the **Bookmarks** pane, choose the context menu (the three vertical dots) for the bookmark that you want to rename, and then choose **rename**.
3. In the **Rename bookmark** pane, enter a name for the bookmark, and then choose **Save**.

Making a bookmark the default view in Amazon Quick Sight

By default, when you update a dashboard, Quick Sight remembers those changes and keeps them after you close the dashboard. This way, you can pick up where you left off when you open the dashboard again. You can set a bookmark as the default view of a dashboard instead. If you do, anytime that you open the dashboard, the bookmark view is presented to you, regardless of the changes you made during your last session.

To set a bookmark as your default view of the dashboard

1. In a published dashboard, choose the bookmark icon at upper right to open the **Bookmarks** pane.
2. In the **Bookmarks** pane, choose the context menu (the three dots) for the bookmark that you want to set as your default view, and then choose **Set as default**.

Sharing bookmarks in Amazon Quick Sight

After you create a bookmark, you can share a URL link for the view with others who have permission to view the dashboard. They can then save that view as their own bookmark.

To share a bookmark with another dashboard subscriber

1. In a published dashboard, choose the bookmark icon at upper right to open the **Bookmarks** pane.
2. In the **Bookmarks** pane, choose the bookmark that you want to share so that the dashboard updates to that view.
3. Choose the share icon at upper right, and then choose **Share this view**.

You can copy the URL link that Quick Sight provides and paste it in an email or IM message to share it with others. The recipient of the URL link can then save the view as their own bookmark. For more information about sharing views of a dashboard, see [Sharing your view of a Amazon Quick Sight dashboard](#).

Deleting bookmarks in Amazon Quick Sight

Use the following procedure to delete a bookmark.

To delete a bookmark

1. In a published dashboard, choose the bookmark icon at upper right to open the **Bookmarks** pane.
2. In the **Bookmarks** pane, choose the context menu (the three vertical dots) for the bookmark that you want to delete, and then choose **Delete**.
3. In the **Delete Bookmark** pane that opens, choose **Yes, Delete Bookmark**.

Gaining insights with machine learning (ML) in Amazon Quick Sight

Amazon Quick Sight uses machine learning to help you uncover hidden insights and trends in your data, identify key drivers, and forecast business metrics. You can also consume these insights in natural language narratives embedded in dashboards.

Using machine learning (ML) and natural language capabilities, Amazon Quick Sight Enterprise Edition takes you beyond descriptive and diagnostic analysis, and launches you into forecasting and decision-making. You can understand your data at a glance, share your findings, and discover the best decisions to achieve your goals. You can do this without developing teams and technology to create the necessary machine learning models and algorithms.

You likely have already built visualizations that answer questions about what happened, when, where, and provide drill down for investigation and identification of patterns. With ML insights, you can avoid spending hours manually analyzing and investigating. You can select from a list of customized context-sensitive narratives, called *autonarratives*, and add them to your analysis. In addition to choosing autonarratives, you can choose to view forecasts, anomalies, and factors contributing to these. You can also add autonarratives that explain the key takeaways in plain language, providing a single data-driven truth for your company.

As time passes and data flows through the system, Amazon Quick Sight continually learns so it can deliver ever more pertinent insights. Instead of deciding what the data means, you can decide what to do with the information it provides.

With a shared foundation based on machine learning, all of your analysts and stakeholders can see trends, anomalies, forecasts, and custom narratives built on millions of metrics. They can see root causes, consider forecasts, evaluate risks, and make well-informed, justifiable decisions.

You can create a dashboard like this with no manual analysis, no custom development skills, and no understanding of machine learning modeling or algorithms. All this capability is built into Amazon Quick Sight Enterprise Edition.

Note

Machine learning capabilities are used as needed throughout the product. Features that actively use machine learning are labeled as such.

With ML Insights, Amazon Quick Sight provides three major features:

- **ML-powered anomaly detection** – Amazon Quick Sight uses Amazon's proven machine learning technology to continuously analyze all your data to detect anomalies (outliers). You can identify the top drivers that contribute to any significant change in your business metrics, such as higher-than-expected sales or a dip in your website traffic. Amazon Quick Sight uses the Random Cut Forest algorithm on millions of metrics and billions of data points. Doing this enables you to get deep insights that are often buried in the aggregates, inaccessible through manual analysis.
- **ML-powered forecasting** – Amazon Quick Sight enables nontechnical users to confidently forecast their key business metrics. The built-in ML Random Cut Forest algorithm automatically handles complex real-world scenarios such as detecting seasonality and trends, excluding outliers, and imputing missing values. You can interact with the data with point-and-click simplicity.
- **Autonarratives** – By using automatic narratives in Amazon Quick Sight, you can build rich dashboards with embedded narratives to tell the story of your data in plain language. Doing this can save hours of sifting through charts and tables to extract the key insights for reporting. It also creates a shared understanding of the data within your organization so you make decisions faster. You can use the suggested autonarrative, or you can customize the computations and language to meet your unique requirements. Amazon Quick Sight is like providing a personal data analyst to all of your users.

Topics

- [Understanding the ML algorithm used by Amazon Quick Sight](#)
- [Dataset requirements for using ML insights with Amazon Quick Sight](#)
- [Working with insights in Amazon Quick Sight](#)
- [Creating autonarratives with Amazon Quick Sight](#)
- [Detecting outliers with ML-powered anomaly detection](#)
- [Forecasting and creating what-if scenarios with Amazon Quick Sight](#)

Understanding the ML algorithm used by Amazon Quick Sight

You don't need any technical experience in machine learning to use the ML-powered features in Amazon Quick Sight. This section dives into the

technical aspects of the algorithm, for those who want the details about how it works. This information isn't required reading to use the features.

Amazon Quick Sight uses a built-in version of the Random Cut Forest (RCF) algorithm. The following sections explain what that means and how it is used in Amazon Quick Sight.

First, let's look at some of the terminology involved:

- **Anomaly** – Something that is characterized by its difference from the majority of the other things in the same sample. Also known as an outlier, an exception, a deviation, and so on.
- **Data point** – A discrete unit—or simply put, a row—in a dataset. However, a row can have multiple data points if you use a measure over different dimensions.
- **Decision Tree** – A way of visualizing the decision process of the algorithm that evaluates patterns in the data.
- **Forecast** – A prediction of future behavior based on current and past behavior.
- **Model** – A mathematical representation of the algorithm or what the algorithm learns.
- **Seasonality** – The repeating patterns of behavior that occur cyclically in time series data.
- **Time series** – An ordered set of date or time data in one field or column.

Topics

- [What's the difference between anomaly detection and forecasting?](#)
- [What is RCF?](#)
- [How RCF is applied to detect anomalies](#)
- [How RCF is applied to generate forecasts](#)
- [References for machine learning and RCF](#)

What's the difference between anomaly detection and forecasting?

Anomaly detection identifies outliers and their contributing drivers to answer the question "What happened that doesn't usually happen?" Forecasting answers the question "If everything continues to happen as expected, what happens in the future?" The math that allows forecasting also enables us to ask "If a few things change, what happens then?"

Both anomaly detection and forecasting begin by examining the current known data points. Amazon Quick Sight anomaly detection begins with what is known so it can establish what is outside the known set, and identify those data points as anomalous (outliers). Amazon Quick Sight forecasting excludes the anomalous data points, and sticks with the known pattern. Forecasting focuses on the established pattern of data distribution. In contrast, anomaly detection focuses on the data points that deviate from what is expected. Each method approaches decision-making from a different direction.

What is RCF?

A *random cut forest* (RCF) is a special type of *random forest* (RF) algorithm, a widely used and successful technique in machine learning. It takes a set of random data points, cuts them down to the same number of points, and then builds a collection of models. In contrast, a model corresponds to a decision tree—thus the name forest. Because RFs can't be easily updated in an incremental manner, RCFs were invented with variables in tree construction that were designed to allow incremental updates.

As an unsupervised algorithm, RCF uses cluster analysis to detect spikes in time series data, breaks in periodicity or seasonality, and data point exceptions. Random cut forests can work as a synopsis or sketch of a dynamic data stream (or a time-indexed sequence of numbers). The answers to our questions about the stream come out of that synopsis. The following characteristics address the stream and how we make connections to anomaly detection and forecasting:

- A *streaming algorithm* is an online algorithm with a small memory footprint. An online algorithm makes its decision about the input point indexed by time t before it sees the $(t+1)$ -st point. The small memory allows nimble algorithms that can produce answers with low latency and allow a user to interact with the data.
- Respecting the ordering imposed by time, as in an *online* algorithm, is necessary in anomaly detection and forecasting. If we already know what will happen the day after tomorrow, then predicting what happens tomorrow isn't a forecast—it's just interpolating an unknown missing value. Similarly, a new product introduced today can be an anomaly, but it doesn't necessarily remain an anomaly at the end of the next quarter.

How RCF is applied to detect anomalies

A human can easily distinguish a data point that stands out from the rest of the data. RCF does the same thing by building a "forest" of decision trees, and then monitoring how new data points change the forest.

An *anomaly* is a data point that draws your attention away from normal points—think of an image of a red flower in a field of yellow flowers. This "displacement of attention" is encoded in the (expected) position of a tree (that is, a model in RCF) that would be occupied by the input point. The idea is to create a forest where each decision tree grows out of a partition of the data sampled for training the algorithm. In more technical terms, each tree builds a specific type of binary space partitioning tree on the samples. As Amazon Quick Sight samples the data, RCF assigns each data point an anomaly score. It gives higher scores to data points that look anomalous. The score is, in approximation, inversely proportional to the resulting depth of the point in the tree. The random cut forest assigns an anomaly score by computing the average score from each constituent tree and scaling the result with respect to the sample size.

The votes or scores of the different models are aggregated because each of the models by itself is a weak predictor. Amazon Quick Sight identifies a data point as anomalous when its score is significantly different from the recent points. What qualifies as an anomaly depends on the application.

The paper [Random Cut Forest Based Anomaly Detection On Streams](#) provides multiple examples of this state-of-the-art online anomaly detection (time-series anomaly detection). RCFs are used on contiguous segments or "shingles" of data, where the data in the immediate segment acts as a context for the most recent one. Previous versions of RCF-based anomaly-detection algorithms score an entire shingle. The algorithm in Amazon Quick Sight also provides an approximate location of the anomaly in the current extended context. This approximate location can be useful in the scenario where there is delay in detecting the anomaly. Delays occur because any algorithm needs to characterize "previously seen deviations" to "anomalous deviations," which can unfold over some time.

How RCF is applied to generate forecasts

To forecast the next value in a stationary time sequence, the RCF algorithm answers the question "What would be the most likely completion, after we have a candidate value?" It uses a single tree in RCF to perform a search for the best candidate. The candidates across different trees are aggregated, because each tree by itself a weak predictor. The aggregation also allows the generation of quantile errors. This process is repeated t times to predict the t -th value in the future.

The algorithm in Amazon Quick Sight is called *BIFOCAL*. It uses two RCFs to create a CALibrated BI-FOrest architecture. The first RCF is used to filter out anomalies and provide a weak forecast, which is corrected by the second. Overall, this approach provides significantly more robust forecasts in comparison to other widely available algorithms such as ETS.

The number of parameters in the Amazon Quick Sight forecasting algorithm is significantly fewer than for other widely available algorithms. This allows it to be useful out of the box, without human adjustment for a larger number of time series data points. As more data accumulates in a particular time series, the forecasts in Amazon Quick Sight can adjust to data drifts and changes of pattern. For time series that show trends, trend detection is performed first to make the series stationary. The forecast of that stationary sequence is projected back with the trend.

Because the algorithm relies on an efficient online algorithm (RCF), it can support interactive "what-if" queries. In these, some of the forecasts can be altered and treated as hypotheticals to provide conditional forecasts. This is the origin of the ability to explore "what-if" scenarios during analysis.

References for machine learning and RCF

To learn more about machine learning and this algorithm, we suggest the following resources:

- The article [Robust Random Cut Forest \(RRCF\): A No Math Explanation](#) provides a lucid explanation without the mathematical equations.
- The book [The Elements of Statistical Learning: Data Mining, Inference, and Prediction, Second Edition \(Springer Series in Statistics\)](#) provides a thorough foundation on machine learning.
- [Random Cut Forest Based Anomaly Detection On Streams](#), a scholarly paper that dives deep into the technicalities of both anomaly detection and forecasting, with examples.

A different approach to RCF appears in other AWS services. If you want to explore how RCF is used in other services, see the following:

- *Amazon Managed Service for Apache Flink SQL Reference*: [RANDOM_CUT_FOREST](#) and [RANDOM_CUT_FOREST_WITH_EXPLANATION](#)
- *Amazon SageMaker Developer Guide*: [Random Cut Forest \(RCF\) Algorithm](#). This approach is also explained in [The Random Cut Forest Algorithm](#), a chapter in [Machine Learning for Business](#) (October 2018).

Dataset requirements for using ML insights with Amazon Quick Sight

To begin using the machine learning capabilities of Amazon Quick Sight, you need to connect to or import your data. You can use an existing Amazon Quick Sight dataset or create a new one. You can directly query your SQL-compatible source, or ingest the data into SPICE.

The data must have the following properties:

- At least one metric (for example, sales, orders, shipped units, sign ups, and so on).
- At least one category dimension (for example, product category, channel, segment, industry, and so on). Categories with NULL values are ignored.
- Anomaly detection requires a minimum of 15 data points for training. For example, if the grain of your data is daily, you need at least 15 days of data. If the grain is monthly, you need at least 15 months of data.
- Forecasting work best with more data. Make sure that your dataset has enough historical data for optimal results. For example, if the grain of your data is daily, you need at least 38 days of data. If the grain is monthly, you need at least 43 months of data. Following are the requirements for each time grain:
 - Years: 32 data points
 - Quarters: 35 data points
 - Months: 43 data points
 - Weeks: 35 data points
 - Days: 38 data points
 - Hours: 39 data points
 - Minutes: 46 data points
 - Seconds: 46 data points
- If you want to analyze anomalies or forecasts, you also need at least one date dimension.

If you don't have a dataset to get started, you can download this sample dataset: [ML Insights Sample Dataset VI](#). After you have a dataset ready, create a new analysis from the dataset.

Working with insights in Amazon Quick Sight

In Amazon Quick Sight, you can add ready-to-use analytical computations to your analysis as widgets. You can work with insights in two ways:

- **Suggested insights**

Amazon Quick Sight creates a list of suggested insights based on its interpretation of the data you put into your visuals. The list changes based on context. In other words, you can see different suggestions depending on what fields you add to your visual and what type of visual you choose.

For example, if you have a time-series visualization, your insights might include period-over-period changes, anomalies, and forecasts. As you add more visualizations to your analysis, you generate more suggested insights.

- **Custom insights**

Custom insights enable you to create your own computation, using your own words to give context to the fields that appear in the widget. When you create a custom insight, you add it to the analysis, and then choose what type of calculation that you want to use. Then, you can add text and formatting to make it look how you want. You can also add more fields, calculations, and parameters.

You can add any combination of suggested and custom insights to your analysis, to create the decision-making environment that best serves your purposes.

Topics

- [Adding suggested insights](#)
- [Adding custom insights to your analysis](#)

Adding suggested insights

Use the following procedure to add suggested insights to your analysis.

Before you begin, make sure that your dataset meets the criteria outlined in [Dataset requirements for using ML insights with Amazon Quick Sight](#).

1. Begin with an analysis that has a few fields added to a visual.
2. On the left, choose **Insights**. The **Insights** panel opens and displays a list of ready-to-use suggested insights.

Each visual also displays a small box on its top border to display how many insights are available for that visual. You can choose this box to open the **Insights** panel, and it opens to whatever view you most recently had open.

Scroll down to preview more insights.

The insights that appear are controlled by the data type of the fields you choose to include in your visual. This list is generated each time you change your visual. If you make changes,

check **Insights** to see what is new. To get a specific insight, see [Adding custom insights to your analysis](#).

3. (Optional) Open the context menu with more options for one of the insights. To do this, choose the ellipses on the top right of the insight (...).

The options are different for each type of insight. The options that you can interact with include the following:

- **Change the time series aggregation** – To year, quarter, month, week, day, hour, or minute.
 - **Analyze contributions to metrics** – Choose contributors and a time frame to analyze.
 - **Show all anomalies** – Browse anomalies in this time frame.
 - **Edit forecast** – Choose forecast length, prediction interval, and seasonality.
 - **Focus on** or **Exclude** – Zoom in or zoom out on your dimensional data.
 - **Show details** – View more information about a recent anomaly (outlier).
 - Provide feedback on the usefulness of the insight in your analysis.
4. Add a suggested insight to your analysis by choosing the plus sign (+) near the insight title.
 5. (Optional) After you add an insight to your analysis, customize the narrative that you want it to display. To do this, choose the v-shaped on-visual menu, then choose **Customize narrative**. For more information, see [Creating autonarratives with Amazon Quick Sight](#).

If your insight is for anomalies (outliers), you can also change the settings for the anomaly detection job. To do this, choose **Configure anomaly**. For more information, see [Setting up ML-powered anomaly detection for outlier analysis](#).

6. (Optional) To remove the insight from your analysis, choose the v-shaped on-visual menu at the top right of the visual. Then choose **Delete**.

Adding custom insights to your analysis

If you don't want to use any of the suggested insights, you can create your own custom insight. Use the following procedure to create a custom computational insight.

1. Start with an existing analysis. On the top menu bar, choose **Add+**. Then choose **Add Insight**.

A container for the new insight is added to the analysis.

2. Do one of the following:

- Choose the computation that you want to use from the list. As you choose each item, an example of that insight's output displays. When you find the one that you want to use, choose **Select**.
- Exit this screen and customize the insight manually. An unconfigured insight has a **Customize insight** button. Choose the button to open the **Configure narrative** screen. For more information on using the expression editor, see [Creating autonarratives with Amazon Quick Sight](#).

Because you are initiating the creation of the insight, it's not based on an existing visual. When the insight is added to the analysis, it displays a note showing what kind of data it needs to complete your request. For example, it might ask for **1 dimension in Time**. In this case, you add a dimension to the **Time** field well.

3. After you have the correct data, follow any remaining screen prompts to finish creating the custom insight.
4. (Optional) To remove the insight from your analysis, choose the v-shaped on-visual menu at the top right of the visual. Then choose **Delete**.

Creating autonarratives with Amazon Quick Sight

An *autonarrative* is a natural-language summary widget that displays descriptive text instead of charts. You can embed these widgets throughout your analysis to highlight key insights and callouts. You don't have to sift through the visual, drilling down, comparing values, and rechecking ideas to extract a conclusion. You also don't have to try to understand what the data means, or discuss different interpretations with your colleagues. Instead, you can extrapolate the conclusion from the data, and display it in the analysis, stated plainly. A single interpretation can be shared by everyone.

Amazon Quick Sight automatically interprets the charts and tables in your dashboard and provides a number of suggested insights in natural language. The suggested insights that you can choose from are ready-made and come with words, calculations, and functions. But you can change them if you want to. You can also design your own. As the author of the dashboard, you have complete flexibility to customize the computations and language for your needs. You can use narratives to effectively tell the story of your data in plain language.

Note

Narratives are separate from machine learning. They only use ML if you add forecast or anomaly (outlier) computations to them.

Topics

- [Insights that include autonarratives](#)
- [Use the narrative expression editor](#)
- [The expression editor workspace](#)
- [Adding URLs](#)
- [Working with autonarrative computations](#)

Insights that include autonarratives

When you are adding an insight, also known as an autonarrative, to your analysis, you can choose from the following templates. In the following list, they are defined by example. Each definition includes a list of the minimum required fields for the autonarrative to work. If you are using only the suggested insights on the **Insights** tab, choose the appropriate fields to get an insight to show up in the suggested insights list.

For more information on customizing autonarratives, see [Working with autonarrative computations](#).

- **Bottom ranked** – For example, the bottom three states by sales revenue. Requires that you have at least one dimension in the **Categories** field well.
- **Bottom movers** – For example, the bottom three products sold, by sales revenue. Requires that you have at least one dimension in the **Time** field well and at least one dimension in the **Categories** field well.
- **Forecast (ML-powered insight)** – For example, "Total sales are forecasted to be \$58,613 for Jan 2016." Requires that you have at least one dimension in the **Time** field well.
- **Growth rate** – For example, "The 3-month compounded growth rate for sales is 22.23%." Requires that you have at least one dimension in the **Time** field well.
- **Maximum** – For example, "Highest month is Nov 2014 with sales of \$112,326." Requires that you have at least one dimension in the **Time** field well.

- **Metric comparison** – For example, "Total sales for Dec 2014 is \$90,474, 10% higher than target of \$81,426." Requires that you have at least one dimension in the **Time** field well and at least two measures in the **Values** field well.
- **Minimum** – For example, "Lowest month is Feb 2011 with sales of \$4,810." Requires that you have at least one dimension in the **Time** field well.
- **Anomaly detection (ML-powered insight)** – For example, top three outliers and their contributing drivers for total sales on January 3, 2019. Requires that you have at least one dimension in the **Time** field well, at least one measure in the **Values** field well, and at least one dimension in the **Categories** field well.
- **Period over period** – For example, "Total sales for Nov 2014 increased by 44.39% (\$34,532) from \$77,793 to \$112,326." Requires that you have at least one dimension in the **Time** field well.
- **Period to date** – For example, "Year-to-date sales for Nov 30, 2014 increased by 25.87% (\$132,236) from \$511,236 to \$643,472." Requires that you have at least one dimension in the **Time** field well.
- **Top ranked** – For example, top three states by sales revenue. Requires that you have at least one dimension in the **Categories** field well.
- **Top movers** – For example, top products by sales revenue for November 2014. Requires that you have at least one dimension in the **Time** field well and at least one dimension in the **Categories** field well.
- **Total aggregation** – For example, "Total revenue is \$2,297,200." Requires that you have at least one dimension in the **Time** field well and at least one measure in the **Values** field well.
- **Unique values** – For example, "There are 793 unique values in Customer_IDs." Requires that you have at least one dimension in the **Categories** field well.

Use the narrative expression editor

The following walkthrough shows an example of how to customize a narrative. For this example, we use a period over period computation type.

1. Begin with an existing analysis. Add a **period over period** insight to it. The easiest way to do this is to choose the + icon, then **Add insight**, then choose a type of insight from the list. To learn what type of computational insights you can add as autonarratives, see [Insights that include autonarratives](#).

After you choose a type of insight, choose **Select** to create the widget. To create an empty narrative, close this screen without choosing a template. To follow this example, choose **Period over period**.

If you had a visual selected when you added the insight, the field wells have preconfigured fields for the date, metric, and category. These come from the visualization that you chose when you created the insight. You can customize the fields as needed.

You can only customize a narrative for a new or existing insight (text-based) widget. You can't add one to an existing visual (chart based), because it's a different type of widget.

2. Edit the narrative in the expressions editor by choosing the on-visual menu, then choosing **Customize narrative**.

In this context, **Computations** are predefined calculations (period-over-period, period-to-date, growth rate, max, min, top movers, and so on) that you can reference in your template to describe your data. Currently, Amazon Quick Sight supports 13 different types of computations that you can add to your insight. In this example, **PeriodOverPeriod** is added by default because we chose the **Period Over Period** template from the suggested insights panel.

3. Choose **Add computation** at bottom right to add a new computation, and then choose one from the list. For this walkthrough, choose **Growth rate**, and then choose **Next**.
4. Configure the computation by choosing the number of periods that you want to compute over. The default is four, and that works for our example. Optionally, you can change the name of the computation at the top of the screen. However, for our purposes, leave the name unchanged.

Note


The computation names that you create are unique within the insight. You can reference multiple computations of the same type in your narrative template. For example, suppose that you have two metrics, sales revenue and units sold. You can create growth rate computations for each metric if they have different names. However, anomaly computations aren't compatible with any other computation type in the same widget. Anomaly detection must exist in an insight by itself. To use other computations in the same analysis, put them into insights separate from anomalies.

To proceed, choose **Add**.

- Expand **Computations** on the right. The computations that are part of the narrative display in the list. In this case, it's **PeriodOverPeriod** and **GrowthRate**.
- In the workspace, add the following text after the final period: **Compounded growth rate for the last**, then add a space.
- Next, to add the computation leave your cursor after the space after the word **last**. On the right, under **GrowthRate**, choose the expression named **timePeriods** (click only once to add it).

Doing this inserts the expression **GrowthRate.timePeriods**, which is the number of periods you set in the configuration for **GrowthRate**.

- Complete the sentence with **days is** (a space before and afterwards), and add the expression **GrowthRate.compoundedGrowthRate.formattedValue**, followed by a period (.). Choose the expression from the list, rather than typing it in. However, you can edit the contents of the expression after you add it.



Note

The **formattedValue** expression returns a string that is formatted based on the formatting applied for the metric on the field. To perform metric math, use **value** instead, which returns the raw value as an integer or decimal.

- Add a conditional statement and formatting. Place your cursor at the end of the template, after the **formattedValue** expression. Add a space if necessary. On the **Edit narrative** menu bar, choose **Insert code**, and then choose **Inline IF** from the list. An expression block opens.

10. With the expression block open, choose **GrowthRate**, **compoundedGrowthRate**, **value** from the expression list. Enter **>0** at the end of the expression. Choose **Save**. Don't move your cursor yet.

A prompt appears for the conditional content; enter **better than expected!** Then select the text you just entered, and use the formatting toolbar at the top to turn it green and bold.
11. Add another expression block for the case when the growth rate wasn't that great by repeating the previous step. But this time, make it **<0** and enter the text **worse than expected**. Make it red instead of green.
12. Choose **Save**. The customized narrative that we just created should look similar to the following.



The expression editor provides you with a sophisticated tool to customize your narratives. You can also reference the parameters you create for your analysis or dashboard, and use a set of built-in functions for further customization.

Tip

To create an empty narrative, add an insight using the **+** icon and then **Add insights**. But instead of choosing a template, simply close the screen.
The best way to get started with customizing narratives is to use the existing templates to learn the syntax.

The expression editor workspace

Use the expression editor to customize a narrative to best fit your business needs. The information below provides an overview of the expression editor workspace and lists all menu options that can be configured for your narrative. For a walkthrough that shows you how to create a custom narrative, see [Use the narrative expression editor](#).

On the right side of the screen, there's a list of items that you can add to the narrative:

- **Computations** – Use this to choose from the computations that are available in this insight. You can expand this list.
- **Parameters** – Use this to choose from the parameters that exist in your analysis. You can expand this list.
- **Functions** – Use this to choose from functions that you can add to a narrative. You can expand this list.
- **Add computation** – Use this button to create another computation. New computations appear in the **Computations** list, ready to add to the insight.

At the bottom of the narrative expression editor, there's a preview of the narrative that updates as you work. This area also shows an alert if you introduce an error into the narrative or if the narrative is empty. To see a preview of ML-powered insights like anomaly detection or forecasting, run your insight calculation at least once before customizing the narrative.

Editing tools are located across the top of the screen. They offer the following options:

- **Insert code** – You can insert the following code blocks from this menu:
 - **Expressions** – Add a free-form expression.
 - **Inline IF** – Add an IF statement that displays inline with the existing block of text.
 - **Inline FOR** – Add a FOR statement that displays inline with the existing block of text.
 - **Block IF** – Add an IF statement that displays in a separate block of text.
 - **Block FOR** – Add a FOR statement that displays in a separate block of text.

The IF and FOR statements enable you to create content that is conditionally formatted. For example, you might add a **block IF** statement, then configure it to compare an integer to a value from a calculation. To do this, you use the following steps, also demonstrated in [Use the narrative expression editor](#):

1. Open the calculations menu at right, and choose one of the blue highlighted items from one of the calculations. Doing this adds the item to the narrative.
2. Click once on the item to open it.
3. Enter the comparison that you want to make. The expression looks something like this:
`PeriodOverPeriod.currentMetricValue.value>0.`
4. Save this expression in the pop-up editor, which prompts you for **Conditional content**.
5. Enter what you want to display in the insight, and format it as you want it to appear. Or if you prefer, you can add an image or a URL—or add a URL to an image.

- **Paragraph** – This menu offers options for changes to the font size:
 - **H1 Large header**
 - H2 Header
 - H3 Small header
 - ¶1 Large paragraph
 - ¶2 Paragraph
 - ¶3 Small paragraph
- **Font** – Use this menu tray to choose options for text formatting. These include bold, italic, underline, strikethrough, foreground color of the text (the letters themselves), and background color of the text. Choose the icon to turn on an option; choose it again to toggle the option off.
- **Formatting** – Use this menu tray to choose options for paragraph formatting, including bulleted list, left justify, center, and right justify. Choose the icon to turn on an option, choose it again to toggle the option off.
- **Image** – Use this icon add an image URL. The image displays in your insight, provided the link is accessible. You can resize images. To display an image based on a condition, put the image inside an IF block.
- **URL** – Use this icon to add a static or dynamic URL. You can also add URLs to images. For example, you can add traffic light indicator images to an insight for an executive dashboard, with links to a new sheet for red, amber, and green conditions.

Adding URLs

Using the **URL** button on the editing menu of the narrative expression editor, you can add static and dynamic URLs (hyperlinks) into a narrative. You can also use the following keyboard shortcuts: ⌘+⬆+L or Ctrl+⬆+L.

A static URL is a link that doesn't change; it always opens the same URL. A dynamic URL is a link that changes based on the expressions or parameters that you provide when you set it up. It's built with dynamically evaluated expressions or parameters.

Following are of examples of when you might add a static link in your narrative:

- **In an IF statement, you might use the URL in the conditional content.** If you do and a metric fails to meet an expected value, your link might send the user to a wiki with a list of best practices to improve the metric.

- **You might use a static URL to create a link to another sheet in the same dashboard, by using the following steps:**

1. Go to the sheet that you want to make the link to.
2. Copy that sheet's URL.
3. Return to the narrative editor and create a link using the URL that you just copied.

Following are examples of when you might add a dynamic link in your narrative:

- **To search a website with a query, by using the following steps.**

1. Create a URL with the following link.

```
https://google.com?q=<<formatDate(now(), 'yyyy-MM-dd')>>
```

This link sends a query to Google with search text that is the evaluated value of the following.

```
formatDate(now(), 'yyyy-MM-dd')
```

If the value of `now()` is 02/02/2020, then the link on your narrative contains `https://google.com?q=2020-02-02`.

- **To create a link that updates a parameter.** To do this, create or edit a link and set the URL to the current dashboard or analysis URL. Then add the expression that sets the parameter value to at the end, for example `#p.myParameter=12345`.

Suppose that the following is the dashboard link that you start with.

```
https://us-east-1.quicksight.aws.amazon.com/sn/analyses/00000000-1111-2222-3333-44444444
```

If you add a parameter value assignment to it, it looks like the following.

```
https://us-east-1.quicksight.aws.amazon.com/sn/analyses/00000000-1111-2222-3333-44444444#p.myParameter=12345
```

For more information on parameters in URLs, see [Using parameters in a URL](#).

Working with autonarrative computations

Use this section to help you understand what functions are available to you when you are customizing an autonarrative. You only need to customize a narrative if you want to change or build on the default computation.

After you create an autonarrative, the expression editor opens. You can also activate the expression editor by choosing the on-visual menu, and then **Customize Narrative**. To add a computation while using the expression editor, choose **+ Add computation**.

You can use the following code expression to build your autonarrative. These are available from the list that's labeled **Insert code**. Code statements can display inline (in a sentence) or as a block (in a list).

- Expression – Create your own code expression.
- IF – An IF statement that includes an expression after evaluating a condition.
- FOR – A FOR statement that loops through values.

You can use the following computations to build your autonarrative. You can use the expression editor without editing any syntax, but you can also customize it if you want to. To interact with the syntax, open the computational widget in the autonarrative expression editor.

Topics

- [ML-powered anomaly detection for outliers](#)
- [Bottom movers computation](#)
- [Bottom ranked computation](#)
- [ML-powered forecasting](#)
- [Growth rate computation](#)
- [Maximum computation](#)
- [Metric comparison computation](#)
- [Minimum computation](#)
- [Period over period computation](#)
- [Period to date computation](#)
- [Top movers computation](#)
- [Top ranked computation](#)

- [Total aggregation computation](#)
- [Unique values computation](#)

ML-powered anomaly detection for outliers

The ML-powered anomaly detection computation searches your data for outliers. For example, you can detect the top three outliers for total sales on January 3, 2019. If you enable contribution analysis, you can also detect the key drivers for each outlier.

To use this function, you need at least one dimension in the **Time** field well, at least one measure in the **Values** field well, and at least one dimension in the **Categories** field well. The configuration screen provides an option to analyze the contribution of other fields as key drivers, even if those fields aren't in the field wells.

For more information, see [Detecting outliers with ML-powered anomaly detection](#).

Note

You can't add ML-powered anomaly detection to another computation, and you can't add another computation to an anomaly detection.

Computation outputs

Each function generates a set of output parameters. You can add these outputs to the autonarrative to customize what it displays. You can also add your own custom text.

To locate the output parameters, open the **Computations** tab on the right, and locate the computation that you want to use. The names of the computations come from the name that you provide when you create the insight. Choose the output parameter by clicking on it only once. If you click twice, you add the same output twice. You can use items displayed in **bold monospace font** following in the narrative.

- `timeField` – From the **Time** field well.
 - `name` – The formatted display name of the field.
 - `timeGranularity` – The time field granularity (**DAY**, **YEAR**, and so on).
- `categoryFields` – From the **Categories** field well.
 - `name` – The formatted display name of the field.

- **metricField** – From the **Values** field well.
 - **name** – The formatted display name of the field.
 - **aggregationFunction** – The aggregation used for the metric (**SUM**, **AVG**, and so on).
- **itemsCount** – The number of items included in this computation.
- **items** – Anomalous items.
 - **timeValue** – The values in the date dimension.
 - **value** – The date/time field at the point of the anomaly (outlier).
 - **formattedValue** – The formatted value in the date/time field at the point of the anomaly.
 - **categoryName** – The actual name of the category (cat1, cat2, and so on).
 - **direction** – The direction on the x-axis or y-axis that's identified as anomalous: HIGH or LOW. HIGH means "higher than expected." LOW means "lower than expected."

When iterating on items, `AnomalyDetection.items[index].direction` can contain either HIGH or LOW. For example, `AnomalyDetection.items[index].direction='HIGH'` or `AnomalyDetection.items[index].direction=LOW`.

`AnomalyDetection.direction` can have an empty string for ALL. An example is `AnomalyDetection.direction=' '`.

- **actualValue** – The metric's actual value at the point of the anomaly or outlier.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the metric field.
 - **formattedAbsoluteValue** – The absolute value formatted by the metric field.
- **expectedValue** – The metric's expected value at the point of the anomaly (outlier).
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the metric field.
 - **formattedAbsoluteValue** – The absolute value formatted by the metric field.

Bottom movers computation

The bottom movers computation counts the requested number of categories by date that rank in the bottom of the autonarrative's dataset. For example, you can create a computation to find the bottom three products sold, by sales revenue.

To use this function, at least one dimension in the **Time** field well and at least one dimension in the **Categories** field well.

Parameters

name

A unique descriptive name that you assign or change. A name is assigned if you don't create your own. You can edit this later.

Date

The date dimension that you want to rank.

Category

The category dimension that you want to rank.

Value

The aggregated measure that the computation is based on.

Number of movers

The number of ranked results that you want to display.

Order by

The order that you want to use, percent difference or absolute difference.

Computation outputs

Each function generates a set of output parameters. You can add these outputs to the autonarrative to customize what it displays. You can also add your own custom text.

To locate the output parameters, open the **Computations** tab on the right, and locate the computation that you want to use. The names of the computations come from the name you provide when you create the insight. Choose the output parameter by clicking on it only once. If you click twice, you add the same output twice. Items displayed in **bold** can be used in the narrative.

Note

These are the same output parameters as the ones that are returned by the top movers computation.

- **timeField** – From the **Time** field well.
 - **name** – The formatted display name of the field.
 - **timeGranularity** – The time field granularity (**DAY**, **YEAR**, and so on).
- **categoryField** – From the **Categories** field well.
 - **name** – The formatted display name of the field.
- **metricField** – From the **Values** field well.
 - **name** – The formatted display name of the field.
 - **aggregationFunction** – The aggregation used for the metric (**SUM**, **AVG**, and so on).
- **startTimeValue** – The value in the date dimension.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the datetime field.
- **endTimeValue** – The value in the date dimension.
 - **value** – The raw value.
 - **formattedValue** – The absolute value formatted by the datetime field.
- **itemsCount** – The number of items included in this computation.
- **items**: Bottom moving items.
 - **categoryField** – The category field.
 - **value** – The value (contents) of the category field.
 - **formattedValue** – The formatted value (contents) of the category field. If the field is null, this displays 'NULL'. If the field is empty, it displays '(empty)'.
 - **currentMetricValue** – The current value for the metric field.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the metric field
 - **formattedAbsoluteValue** – The absolute value formatted by the metric field.
 - **previousMetricValue** – The previous value for the metric field.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the metric field
 - **formattedAbsoluteValue** – The absolute value formatted by the metric field.
 - **percentDifference** – The percent difference between the current and previous values of the metric field.
 - **value** – The raw value of the calculation of the percent difference.

- **formattedValue** – The formatted value of the percent difference (for example, -42%).
- **formattedAbsoluteValue** – The formatted absolute value of the percent difference (for example, 42%).
- **absoluteDifference** – The absolute difference between the current and previous values of the metric field.
 - **value** – The raw value of the calculation of the absolute difference.
 - **formattedValue** – The absolute difference formatted by the settings in the metric field's format preferences.
 - **formattedAbsoluteValue** – The absolute value of the difference formatted by the metric field.

Bottom ranked computation

The bottom ranked computation calculates the requested number of categories by value that rank in the bottom of the autonarrative's dataset. For example, you can create a computation to find the bottom three states by sales revenue.

To use this function, you need at least one dimension in the **Categories** field well.

Parameters

name

A unique descriptive name that you assign or change. A name is assigned if you don't create your own. You can edit this later.

Category

The category dimension that you want to rank.

Value

The aggregated measure that the computation is based on.

Number of results

The number of ranked results that you want to display.

Computation outputs

Each function generates a set of output parameters. You can add these outputs to the autonarrative to customize what it displays. You can also add your own custom text.

To locate the output parameters, open the **Computations** tab on the right, and locate the computation that you want to use. The names of the computations come from the name you provide when you create the insight. Choose the output parameter by clicking on it only once. If you click twice, you add the same output twice. Items displayed in **bold** can be used in the narrative.


Note

These are the same output parameters as the ones that are returned by the top ranked computation.

- **categoryField** – From the **Categories** field well.
 - **name** – The formatted display name of the field.
- **metricField** – From the **Values** field well.
 - **name** – The formatted display name of the field.
 - **aggregationFunction** – The aggregation used for the metric (**SUM**, **AVG**, and so on).
- **itemsCount** – The number of items included in this computation.
- **items**: Bottom ranked items.
 - **categoryField** – The category field.
 - **value** – The value (contents) of the category field.
 - **formattedValue** – The formatted value (contents) of the category field. If the field is null, this displays 'NULL'. If the field is empty, it displays '(empty)'.
 - **metricValue** – The metric field.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the metric field.
 - **formattedAbsoluteValue** – The absolute value formatted by the metric field.

Example

The following screenshot shows the default configuration for the bottom-ranked computation.



ML-powered forecasting

The ML-powered forecast computation forecasts future metrics based on patterns of previous metrics by seasonality. For example, you can create a computation to forecast total revenue for the next six months.

To use this function, you need at least one dimension in the **Time** field well.

For more information about working with forecasts, see [Forecasting and creating what-if scenarios with Amazon Quick Sight](#).

Parameters

name

A unique descriptive name that you assign or change. A name is assigned if you don't create your own. You can edit this later.

Date

The date dimension that you want to rank.

Value

The aggregated measure that the computation is based on.

Periods forward

The number of time periods in the future that you want to forecast. Ranges from 1 to 1,000.

Periods backward

The number of time periods in the past that you want to base your forecast on. Ranges from 0 to 1,000.

Seasonality

The number of seasons included in the calendar year. The default setting, **automatic** detects this for you. Ranges from 1 to 180.

Computation outputs

Each function generates a set of output parameters. You can add these outputs to the autonarrative to customize what it displays. You can also add your own custom text.

To locate the output parameters, open the **Computations** tab on the right, and locate the computation that you want to use. The names of the computations come from the name you provide when you create the insight. Choose the output parameter by clicking on it only once. If you click twice, you add the same output twice. Items displayed in **bold** can be used in the narrative.

- `timeField` – From the **Time** field well.
 - **name** – The formatted display name of the field.
 - **timeGranularity** – The time field granularity (**DAY**, **YEAR**, and so on).
- `metricField` – From the **Values** field well.
 - **name** – The formatted display name of the field.
 - **aggregationFunction** – The aggregation used for the metric (**SUM**, **AVG**, and so on).
- `metricValue` – The value in the metric dimension.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the metric field.
 - **formattedAbsoluteValue** – The absolute value formatted by the metric field.
- `timeValue` – The value in the date dimension.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the date field.
- **relativePeriodsToForecast** – The relative number of periods between latest datetime record and last forecast record.

Growth rate computation

The growth rate computation compares values over time periods. For example, you can create a computation to find the three-month compounded growth rate for sales, expressed as a percentage.

To use this function, you need at least one dimension in the **Time** field well.

Parameters

name

A unique descriptive name that you assign or change. A name is assigned if you don't create your own. You can edit this later.

Date

The date dimension that you want to rank.

Value

The aggregated measure that the computation is based on.

Number of periods

The number of time periods in the future that you want to use to compute the growth rate.

Computation outputs

Each function generates a set of output parameters. You can add these outputs to the autonarrative to customize what it displays. You can also add your own custom text.

To locate the output parameters, open the **Computations** tab on the right, and locate the computation that you want to use. The names of the computations come from the name you provide when you create the insight. Choose the output parameter by clicking on it only once. If you click twice, you add the same output twice. Items displayed in **bold** can be used in the narrative.

- `timeField` – From the **Time** field well.
 - **name** – The formatted display name of the field.
 - **timeGranularity** – The time field granularity (**DAY**, **YEAR**, and so on).
- `metricField` – From the **Values** field well.
 - **name** – The formatted display name of the field.

- **aggregationFunction** – The aggregation used for the metric (**SUM**, **AVG**, and so on).
- **previousMetricValue** – The previous value in the metric dimension.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the metric field.
 - **formattedAbsoluteValue** – The absolute value formatted by the metric field.
- **previousTimeValue** – The previous value in the datetime dimension.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the datetime field.
- **compoundedGrowthRate** – The percent difference between the current and previous values of the metric field.
 - **value** – The raw value of the calculation of the percent difference.
 - **formattedValue** – The formatted value of the percent difference (for example, -42%).
 - **formattedAbsoluteValue** – The formatted absolute value of the percent difference (for example, 42%).
- **absoluteDifference** – The absolute difference between the current and previous values of the metric field.
 - **value** – The raw value of the calculation of the absolute difference.
 - **formattedValue** – The absolute difference formatted by the settings in the metric field's format preferences.
 - **formattedAbsoluteValue** – The absolute value of the difference formatted by the metric field.

Maximum computation

The maximum computation finds the maximum dimension by value. For example, you can create a computation to find the month with the highest revenue.

To use this function, you need at least one dimension in the **Time** field well.

Parameters

name

A unique descriptive name that you assign or change. A name is assigned if you don't create your own. You can edit this later.

Date

The date dimension that you want to rank.

Value

The aggregated measure that the computation is based on.

Computation outputs

Each function generates a set of output parameters. You can add these outputs to the autonarrative to customize what it displays. You can also add your own custom text.

To locate the output parameters, open the **Computations** tab on the right, and locate the computation that you want to use. The names of the computations come from the name you provide when you create the insight. Choose the output parameter by clicking on it only once. If you click twice, you add the same output twice. Items displayed in **bold** can be used in the narrative.

Note

These are the same output parameters as the ones that are returned by the minimum computation.

- `timeField` – From the **Time** field well.
 - **name** – The formatted display name of the field.
 - **timeGranularity** – The time field granularity (**DAY**, **YEAR**, and so on).
- `metricField` – From the **Values** field well.
 - **name** – The formatted display name of the field.
 - **aggregationFunction** – The aggregation used for the metric (**SUM**, **AVG**, and so on).
- `metricValue` – The value in the metric dimension.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the metric field.
 - **formattedAbsoluteValue** – The absolute value formatted by the metric field.
- `timeValue` – The value in the datetime dimension.
 - **value** – The raw value.

- **formattedValue** – The value formatted by the datetime field.

Metric comparison computation

The metric comparison computation compares values in different measures. For example, you can create a computation to compare two values, such as actual sales compared to sales goals.

To use this function, you need at least one dimension in the **Time** field well and at least two measures in the **Values** field well.

Parameters

name

A unique descriptive name that you assign or change. A name is assigned if you don't create your own. You can edit this later.

Date

The date dimension that you want to rank.

Value

The aggregated measure that the computation is based on.

Target value

The field that you want to compare to the value.

Computation outputs

Each function generates a set of output parameters. You can add these outputs to the autonarrative to customize what it displays. You can also add your own custom text.

To locate the output parameters, open the **Computations** tab on the right, and locate the computation that you want to use. The names of the computations come from the name you provide when you create the insight. Choose the output parameter by clicking on it only once. If you click twice, you add the same output twice. Items displayed in **bold** can be used in the narrative.

- **timeField** – From the **Time** field well.
- **name** – The formatted display name of the field.

- **timeGranularity** – The time field granularity (**DAY**, **YEAR**, and so on).
- **fromMetricField** – From the **Values** field well.
 - **name** – The formatted display name of the field.
 - **aggregationFunction** – The aggregation used for the metric (**SUM**, **AVG**, and so on).
- **fromMetricValue** – The value in the metric dimension.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the metric field.
 - **formattedAbsoluteValue** – The absolute value formatted by the metric field.
- **toMetricField** – From the **Values** field well.
 - **name** – The formatted display name of the field.
 - **aggregationFunction** – The aggregation used for the metric (**SUM**, **AVG**, and so on).
- **toMetricValue** – The current value in the metric dimension.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the metric field.
 - **formattedAbsoluteValue** – The absolute value formatted by the metric field.
- **timeValue** – The value in the datetime dimension.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the datetime field.
- **percentDifference** – The percent difference between the current and previous values of the metric field.
 - **value** – The raw value of the calculation of the percent difference.
 - **formattedValue** – The formatted value of the percent difference (for example, -42%).
 - **formattedAbsoluteValue** – The formatted absolute value of the percent difference (for example, 42%).
- **absoluteDifference** – The absolute difference between the current and previous values of the metric field.
 - **value** – The raw value of the calculation of the absolute difference.
 - **formattedValue** – The absolute difference formatted by the settings in the metric field's format preferences.
 - **formattedAbsoluteValue** – The absolute value of the difference formatted by the metric field.

Minimum computation

The minimum computation finds the minimum dimension by value. For example, you can create a computation to find the month with the lowest revenue.

To use this function, you need at least one dimension in the **Time** field well.

Parameters

name

A unique descriptive name that you assign or change. A name is assigned if you don't create your own. You can edit this later.

Date

The date dimension that you want to rank.

Value

The aggregated measure that the computation is based on.

Computation outputs

Each function generates a set of output parameters. You can add these outputs to the autonarrative to customize what it displays. You can also add your own custom text.

To locate the output parameters, open the **Computations** tab on the right, and locate the computation that you want to use. The names of the computations come from the name you provide when you create the insight. Choose the output parameter by clicking on it only once. If you click twice, you add the same output twice. Items displayed in **bold** can be used in the narrative.

Note

These are the same output parameters as the ones that are returned by the maximum computation.

- `timeField` – From the **Time** field well.
- **name** – The formatted display name of the field.
- **timeGranularity** – The time field granularity (**DAY**, **YEAR**, and so on).

- **metricField** – From the **Values** field well.
 - **name** – The formatted display name of the field.
 - **aggregationFunction** – The aggregation used for the metric (**SUM**, **AVG**, and so on).
- **metricValue** – The value in the metric dimension.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the metric field.
 - **formattedAbsoluteValue** – The absolute value formatted by the metric field.
- **timeValue** – The value in the datetime dimension.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the datetime field.

Period over period computation

The period over period computation compares values from two different time periods. For example, you can create a computation to find out how much sales increased or decreased since the previous time period.

To use this function, you need at least one dimension in the **Time** field well.

Parameters

name

A unique descriptive name that you assign or change. A name is assigned if you don't create your own. You can edit this later.

Date

The date dimension that you want to rank.

Value

The aggregated measure that the computation is based on.

Computation outputs

Each function generates a set of output parameters. You can add these outputs to the autonarrative to customize what it displays. You can also add your own custom text.

To locate the output parameters, open the **Computations** tab on the right, and locate the computation that you want to use. The names of the computations come from the name you provide when you create the insight. Choose the output parameter by clicking on it only once. If you click twice, you add the same output twice. Items displayed in **bold** can be used in the narrative.


- `timeField` – From the **Time** field well.
 - **name** – The formatted display name of the field.
 - **timeGranularity** – The time field granularity (**DAY**, **YEAR**, and so on).
- `metricField` – From the **Values** field well.
 - **name** – The formatted display name of the field.
 - **aggregationFunction** – The aggregation used for the metric (**SUM**, **AVG**, and so on).
- `previousMetricValue` – The previous value in the metric dimension.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the metric field.
 - **formattedAbsoluteValue** – The absolute value formatted by the metric field.
- `previousTimeValue` – The previous value in the datetime dimension.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the datetime field.
- `currentMetricValue` – The current value in the metric dimension.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the metric field.
 - **formattedAbsoluteValue** – The absolute value formatted by the metric field.
- `currentTimeValue` – The current value in the datetime dimension.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the datetime field.
- `percentDifference` – The percent difference between the current and previous values of the metric field.
 - **value** – The raw value of the calculation of the percent difference.
 - **formattedValue** – The formatted value of the percent difference (for example, -42%).
 - **formattedAbsoluteValue** – The formatted absolute value of the percent difference (for example, 42%).

- **absoluteDifference** – The absolute difference between the current and previous values of the metric field.
- **value** – The raw value of the calculation of the absolute difference.
- **formattedValue** – The absolute difference formatted by the settings in the metric field's format preferences.
- **formattedAbsoluteValue** – The absolute value of the difference formatted by the metric field.

Example

To create a Period over period computation

1. In the analysis that you want to change, choose **Add insight**.
2. For **Computation type**, choose **Period over period**, and then choose **Select**.
3. In the new insight that you created, add the time dimension and value dimension fields that you want to compare. In the screenshot below, **Order Date** and **Sales (Sum)** are added to the insight. With these two fields selected, Quick Sight shows the year to date sales of the latest month and the percentage difference compared with the previous month.



4. (Optional) To further customize the insight, open the on-visual menu and choose **Customize narrative**. In the **Edit narrative** window that appears, drag and drop the fields that you need from the **Computations** list, and then choose **Save**.

Period to date computation

The period to date computation evaluates values for a specified period to date. For example, you can create a computation to find out how much you've earned in year-to-date sales.

To use this function, you need at least one dimension in the **Time** field well.

Parameters

name

A unique descriptive name that you assign or change. A name is assigned if you don't create your own. You can edit this later.

Date

The date dimension that you want to rank.

Value

The aggregated measure that the computation is based on.

Time granularity

The date granularity that you want to use for the computation, for example year to date.

Computation outputs

Each function generates a set of output parameters. You can add these outputs to the autonarrative to customize what it displays. You can also add your own custom text.

To locate the output parameters, open the **Computations** tab on the right, and locate the computation that you want to use. The names of the computations come from the name you provide when you create the insight. Choose the output parameter by clicking on it only once. If you click twice, you add the same output twice. Items displayed in **bold** can be used in the narrative.

- `timeField` – From the **Time** field well.
 - **name** – The formatted display name of the field.
 - **timeGranularity** – The time field granularity (**DAY**, **YEAR**, and so on).
- `metricField` – From the **Values** field well.
 - **name** – The formatted display name of the field.
 - **aggregationFunction** – The aggregation used for the metric (**SUM**, **AVG**, and so on).
- `previousMetricValue` – The previous value in the metric dimension.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the metric field.


- **formattedAbsoluteValue** – The absolute value formatted by the metric field.
- **previousTimeValue** – The previous value in the datetime dimension.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the datetime field.
- **currentMetricValue** – The current value in the metric dimension.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the metric field.
 - **formattedAbsoluteValue** – The absolute value formatted by the metric field.
- **currentTimeValue** – The current value in the datetime dimension.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the datetime field.
- **periodGranularity** – The period granularity for this computation (**MONTH**, **YEAR**, and so on).
- **percentDifference** – The percent difference between the current and previous values of the metric field.
 - **value** – The raw value of the calculation of the percent difference.
 - **formattedValue** – The formatted value of the percent difference (for example, -42%).
 - **formattedAbsoluteValue** – The formatted absolute value of the percent difference (for example, 42%).
- **absoluteDifference** – The absolute difference between the current and previous values of the metric field.
 - **value** – The raw value of the calculation of the absolute difference.
 - **formattedValue** – The absolute difference formatted by the settings in the metric field's format preferences.
 - **formattedAbsoluteValue** – The absolute value of the difference formatted by the metric field.

Example

To create a Period to date computation

1. In the analysis that you want to change, choose **Add insight**.
2. For **Computation type**, choose **Period to date**, and then choose **Select**.

3. In the new insight that you created, add the time dimension and value dimension fields that you want to compare. In the screenshot below, **Order Date** and **Sales (Sum)** are added to the insight. With these two fields selected, Quick Sight shows the year to date sales of the latest month and the percentage difference compared with the previous month.



4. (Optional) To further customize the insight, open the on-visual menu and choose **Customize narrative**. In the **Edit narrative** window that appears, drag and drop the fields that you need from the **Computations** list, and then choose **Save**.

Top movers computation

The top movers computation counts the requested number of categories by date that rank in the top of the autonarrative's dataset. For example, you can create a computation to find the top products by sales revenue for a time period.

To use this function, you need at least one dimension in the **Time** field well and at least one dimension in the **Categories** field well.

Parameters

name

A unique descriptive name that you assign or change. A name is assigned if you don't create your own. You can edit this later.

Category

The category dimension you want to rank.

Value

The aggregated measure that the computation is based on.

Number of results

The number of top ranking items you want to find.

Computation outputs

Each function generates a set of output parameters. You can add these outputs to the autonarrative to customize what it displays. You can also add your own custom text.

To locate the output parameters, open the **Computations** tab on the right, and locate the computation that you want to use. The names of the computations come from the name you provide when you create the insight. Choose the output parameter by clicking on it only once. If you click twice, you add the same output twice. Items displayed in **bold** can be used in the narrative.

Note

These are the same output parameters as the ones that are returned by the bottom movers computation.

- `timeField` – From the **Time** field well.
 - **name** – The formatted display name of the field.
 - **timeGranularity** – The time field granularity (**DAY**, **YEAR**, and so on).
- `categoryField` – From the **Categories** field well.
 - **name** – The formatted display name of the field.
- `metricField` – From the **Values** field well.
 - **name** – The formatted display name of the field.
 - **aggregationFunction** – The aggregation used for the metric (**SUM**, **AVG**, and so on).
- `startTimeValue` – The value in the date dimension.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the datetime field.
- `endTimeValue` – The value in the date dimension.
 - **value** – The raw value.
 - **formattedValue** – The absolute value formatted by the datetime field.
- **itemsCount** – The number of items included in this computation.
- **items**: Top moving items.
 - `categoryField` – The category field.

- **value** – The value (contents) of the category field.
- **formattedValue** – The formatted value (contents) of the category field. If the field is null, this displays 'NULL'. If the field is empty, it displays '(empty)'.
- **currentMetricValue** – The current value for the metric field.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the metric field.
 - **formattedAbsoluteValue** – The absolute value formatted by the metric field.
- **previousMetricValue** – The previous value for the metric field.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the metric field.
 - **formattedAbsoluteValue** – The absolute value formatted by the metric field.
- **percentDifference** – The percent difference between the current and previous values of the metric field.
 - **value** – The raw value of the calculation of the percent difference.
 - **formattedValue** – The formatted value of the percent difference (for example, -42%).
 - **formattedAbsoluteValue** – The formatted absolute value of the percent difference (for example, 42%).
- **absoluteDifference** – The absolute difference between the current and previous values of the metric field.
 - **value** – The raw value of the calculation of the absolute difference.
 - **formattedValue** – The absolute difference formatted by the settings in the metric field's format preferences.
 - **formattedAbsoluteValue** – The absolute value of the difference formatted by the metric field.

Top ranked computation

The top ranked computation finds the top ranking dimensions by value. For example, you can create a computation to find the top three states by sales revenue.

To use this function, you need at least one dimension in the **Categories** field well.

Parameters

name

A unique descriptive name that you assign or change. A name is assigned if you don't create your own. You can edit this later.

Category

The category dimension that you want to rank.

Value

The aggregated measure that the computation is based on.

Number of results

The number of top ranking items that you want to find.

Computation outputs

Each function generates a set of output parameters. You can add these outputs to the autonarrative to customize what it displays. You can also add your own custom text.

To locate the output parameters, open the **Computations** tab on the right, and locate the computation that you want to use. The names of the computations come from the name you provide when you create the insight. Choose the output parameter by clicking on it only once. If you click twice, you add the same output twice. Items displayed in **bold** can be used in the narrative.

Note

These are the same output parameters as the ones that are returned by the bottom ranked computation.

- `categoryField` – From the **Categories** field well.
 - **name** – The formatted display name of the field.
- `metricField` – From the **Values** field well.
 - **name** – The formatted display name of the field.
 - **aggregationFunction** – The aggregation used for the metric (**SUM**, **AVG**, and so on).

- **itemsCount** – The number of items included in this computation.
- **items**: Top ranked items.
 - **categoryField** – The category field.
 - **value** – The value (contents) of the category field.
 - **formattedValue** – The formatted value (contents) of the category field. If the field is null, this displays 'NULL'. If the field is empty, it displays '(empty)'.
 - **metricValue** – The metric field.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the metric field.
 - **formattedAbsoluteValue** – The absolute value formatted by the metric field.

Total aggregation computation

The total aggregation computation creates a grand total of the value. For example, you can create a computation to find the total revenue.

To use this function, you need at least one dimension in the **Time** field well and at least one measure in the **Values** field well.

Parameters

name

A unique descriptive name that you assign or change. A name is assigned if you don't create your own. You can edit this later.

Value

The aggregated measure that the computation is based on.

Computation outputs

Each function generates a set of output parameters. You can add these outputs to the autonarrative to customize what it displays. You can also add your own custom text.

To locate the output parameters, open the **Computations** tab on the right, and locate the computation that you want to use. The names of the computations come from the name you provide when you create the insight. Choose the output parameter by clicking on it only once.

If you click twice, you add the same output twice. Items displayed in **bold** can be used in the narrative.

- **categoryField** – The category field.
 - **name** – The display name of the category field.
- **metricField** – From the **Values** field well.
 - **name** – The formatted display name of the field.
 - **aggregationFunction** – The aggregation used for the metric (**SUM**, **AVG**, and so on).
- **totalAggregate** – The total value of the metric aggregation.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the metric field.
 - **formattedAbsoluteValue** – The absolute value formatted by the metric field.

Unique values computation

The unique values computation counts the unique values in a category field. For example, you can create a computation to count the number of unique values in a dimension, such as how many customers you have

To use this function, you need at least one dimension in the **Categories** field well.

Parameters

name

A unique descriptive name that you assign or change. A name is assigned if you don't create your own. You can edit this later.

Category

The category dimension that you want to rank.

Computation outputs

Each function generates a set of output parameters. You can add these outputs to the autonarrative to customize what it displays. You can also add your own custom text.

To locate the output parameters, open the **Computations** tab on the right, and locate the computation that you want to use. The names of the computations come from the name you

provide when you create the insight. Choose the output parameter by clicking on it only once. If you click twice, you add the same output twice. Items displayed in **bold** can be used in the narrative.

- `categoryField` – The category field.
 - **name** – The display name of the category field.
- **uniqueGroupValuesCount** – The number of unique values included in this computation.

Detecting outliers with ML-powered anomaly detection

Amazon Quick Sight uses proven Amazon technology to continuously run ML-powered anomaly detection across millions of metrics to discover hidden trends and outliers in your data. This tool allows you to get deep insights that are often buried in the aggregates and not scalable with manual analysis. With ML-powered anomaly detection, you can find outliers in your data without the need for manual analysis, custom development, or ML domain expertise.

Amazon Quick Sight notifies you in your visuals if it detects that you can analyze an anomaly or do some forecasting on your data.

Anomaly detection is not available in the eu-central-2 Europe (Zurich) region.

Important

ML-powered anomaly detection is a compute-intensive task. Before you start using it, you can get an idea of costs by analyzing the amount of data that you want to use. We offer a tiered pricing model that is based on the number of metrics you process per month.

Topics

- [Concepts for anomaly or outlier detection](#)
- [Setting up ML-powered anomaly detection for outlier analysis](#)
- [Exploring outliers and key drivers with ML-powered anomaly detection and contribution analysis](#)

Concepts for anomaly or outlier detection

Amazon Quick Sight uses the word *anomaly* to describe data points that fall outside an overall pattern of distribution. There are many other words for anomalies, which is a scientific term,

including outliers, deviations, oddities, exceptions, irregularities, quirks, and many more. The term that you use might be based on the type of analysis you do, or the type of data you use, or even just the preference of your group. These outlying data points represent an entity—a person, place, thing, or time—which is exceptional in some way.

Humans easily recognize patterns and spot things that aren't like the others. Our senses provide this information for us. If the pattern is simple, and there is only a little data, you can easily make a graph to highlight the outliers in your data. Some simple examples include the following:

- A red balloon in a group of blue ones
- A racehorse that is far ahead of the others
- A kid who isn't paying attention during class
- A day when online orders are up, but shipping is down
- A person who got well, where others didn't

Some data points represent a significant event, and others represent a random occurrence. Analysis uncovers which data is worth investigating, based on what driving factors (key drivers) contributed to the event. Questions are essential to data analysis. Why did it happen? What's it related to? Did it happen only once or many times? What can you do to encourage or discourage more like it?

Understanding how and why a variation exists, and whether there is a pattern in the variations, requires more thought. Without the assistance of machine learning, each person might come to a different conclusion, because they have different experience and information. Therefore, each person might make a slightly different business decision. If there is a lot of data or variables to consider, it can require an overwhelming amount of analysis.

ML-powered anomaly detection identifies the causations and correlations to enable you to make data-driven decisions. You still have control over defining how you want the job to work on your data. You can specify your own parameters, and choose additional options, such as identifying key drivers in a contribution analysis. Or you can use the default settings. The following section walks you through the setup process, and provides explanations for the options available.

Setting up ML-powered anomaly detection for outlier analysis

Use procedures in the following sections to start detecting outliers, detecting anomalies, and identifying the key drivers that contribute to them.

Topics

- [Viewing anomaly and forecast notifications](#)
- [Adding an ML insight to detect outliers and key drivers](#)
- [Using contribution analysis for key drivers](#)

Viewing anomaly and forecast notifications

Amazon Quick Sight notifies you on a visual where it detects an anomaly, key drivers, or a forecasting opportunity. You can follow the prompts to set up anomaly detection or forecasting based on the data in that visual.

1. In an existing line chart, look for an insight notification in the menu on the visual widget.
2. Choose the lightbulb icon to display the notification.
3. If you want more information about the ML insight, you can follow the screen prompts to add an ML insight.

Adding an ML insight to detect outliers and key drivers

You can add an ML insight that detects *anomalies*, which are outliers that seem significant. To get started, you create for your insight a widget, also known as an *autonarrative*. As you configure your options, you can view a limited screenshot of your insight in the **Preview** pane at screen right.

In your insight widget, you can add up to five dimension fields that are not calculated fields. In the field wells, values for **Categories** represent the dimensional values that Amazon Quick Sight uses to split the metric. For example, let's say that you are analyzing revenue across all product categories and product SKUs. There are 10 product categories, each with 10 product SKUs. Amazon Quick Sight splits the metric by the 100 unique combinations and runs anomaly detection on each combination for the split.

The following procedure shows how to do this, and also how to add contribution analysis to detect the key drivers that are causing each anomaly. You can add contribution analysis later, as described in [Using contribution analysis for key drivers](#).

To set up outlier analysis, including key drivers

1. Open your analysis and in the toolbar, choose **Insights**, then **Add**. From the list, choose **Anomaly detection** and **Select**.
2. Follow the screen prompt on the new widget, which tells you to choose fields for the insight. Add at least one date, one measure, and one dimension.

3. Choose **Get started** on the widget. The configuration screen appears.
4. Under **Compute options**, choose values for the following options:
 - a. For **Combinations to be analysed**, choose one of the following options:

- i. **Hierarchical**

Choose this option if you want to analyze the fields hierarchically. For example, if you chose a date (T), a measure (N), and three dimension categories (C1, C2, and C3), Quick Sight analyses the fields hierarchically, as shown following.

T-N, T-C1-N, T-C1-C2-N, T-C1-C2-C3-N

- ii. **Exact**

Choose this option if you want to analyze only the exact combination of fields in the Category field well, as they are listed. For example, if you chose a date (T), a measure (N), and three dimension categories (C1, C2, and C3), Quick Sight analyses only the exact combination of category fields in the order they are listed, as shown following.

T-C1-C2-C3-N

- iii. **All**

Choose this option if you want to analyze all field combinations in the Category field well. For example, if you chose a date (T), a measure (N), and three dimension categories (C1, C2, and C3), Quick Sight analyses all combinations of fields, as shown following.

T-N, T-C1-N, T-C1-C2-N, T-C1-C2-C3-N, T-C1-C3-N, T-C2-N, T-C2-C3-N, T-C3-N

If you chose a date and a measure only, Quick Sight analyses the fields by date and then by measure.

In the **Fields to be analyzed** section, you can see a list of fields from the field wells for reference.

- b. For **Name**, enter a descriptive alphanumeric name with no spaces, or choose the default value. This provides a name for the computation.

If you plan on editing the narrative that automatically displays on the widget, you can use the name to identify this widget's calculation. Customize the name if you plan to edit the autonarrative and if you have other similar calculations in your analysis.

5. In the **Display options** section, choose the following options to customize what is displayed in your insight widget. You can still explore all your results, no matter what you display.
 - a. **Maximum number of anomalies to show** – The number of outliers you want to display in the narrative widget.
 - b. **Severity** – The minimum level of severity for anomalies that you want to display in the insight widget.

A level of severity is a range of anomaly scores that is characterized by the lowest actual anomaly score included in the range. All anomalies that score higher are included in the range. If you set severity to **Low**, the insight displays all of the anomalies that rank between low and very high. If you set the severity to **Very high**, the insight displays only the anomalies that have the highest anomaly scores.

You can use the following options:

- **Very high**
 - **High and above**
 - **Medium and above**
 - **Low and above**
- c. **Direction** – The direction on the x-axis or y-axis that you want to identify as anomalous. You can choose from the following:
 - **Higher than expected** to identify higher values as anomalies.
 - **Lower than expected** to identify lower values as anomalies.
 - **[ALL]** to identify all anomalous values, high and low (default setting).
 - d. **Delta** – Enter a custom value to use to identify anomalies. Any amount higher than the threshold value counts as an anomaly. The values here change how the insight works in your analysis. In this section, you can set the following:
 - **Absolute value** – The actual value to use. For example, suppose this is 48. Amazon Quick Sight then identifies values as anomalous when the difference between a value and the expected value is greater than 48.

- **Percentage** – The percentage threshold to use. For example, suppose this is 12.5%. Amazon Quick Sight then identifies values as anomalous when the difference between a value and the expected value is greater than 12.5%.
- e. **Sort by** – Choose a sort method for your results. Some methods are based on the anomaly score that Amazon Quick Sight generates. Amazon Quick Sight gives higher scores to data points that look anomalous. You can use any of the following options:
- **Weighted anomaly score** – The anomaly score multiplied by the log of the absolute value of the difference between the actual value and the expected value. This score is always a positive number.
 - **Anomaly score** – The actual anomaly score assigned to this data point.
 - **Weighted difference from expected value** – The anomaly score multiplied by the difference between the actual value and the expected value (default).
 - **Difference from expected value** – The actual difference between the actual value and the expected value (that is, actual–expected).
 - **Actual value** – The actual value with no formula applied.
6. In the **Schedule options** section, set the schedule for automatically running the insight recalculation. The schedule runs only for published dashboards. In the analysis, you can run it manually as needed. Scheduling includes the following settings:
- **Occurrence** – How often that you want the recalculation to run: every hour, every day, every week, or every month.
 - **Start schedule on** – The date and time to start running this schedule.
 - **Timezone** – The time zone that the schedule runs in. To view a list, delete the current entry.
7. In the **Top contributors** section, set Amazon Quick Sight to analyze the key drivers when an outlier (anomaly) is detected.

For example, Amazon Quick Sight can show the top customers that contributed to a spike in sales in the US for home improvement products. You can add up to four dimensions from your dataset. These include dimensions that you didn't add to the field wells of this insight widget.

For a list of dimensions available for contribution analysis, choose **Select fields**.

8. Choose **Save** to confirm your choices. Choose **Cancel** to exit without saving.
9. From the insight widget, choose **Run now** to run the anomaly detection and view your insight.

The amount of time that anomaly detection takes to complete varies depending on how many unique data points you are analyzing. The process can take a few minutes for a minimum number of points, or it can take many hours.

While it's running in the background, you can do other work in your analysis. Make sure to wait for it to complete before you change the configuration, edit the narrative, or open the **Explore anomalies** page for this insight.

The insight widget needs to run at least once before you can see results. If you think the status might be out of date, you can refresh the page. The insight can have the following states.

Appears on the Page	Status
Run now button	The job has not yet started.
Message about Analyzing for anomalies	The job is currently running.
Narrative about the detected anomalies (outliers)	The job has run successfully. The message says when this widget's calculation was last updated.
Alert icon with an exclamation point (!)	This icon indicates there was an error during the last run. If the narrative also displays, you can still use Explore anomalies to use data from the previous successful run.

Using contribution analysis for key drivers

Amazon Quick Sight can identify the dimensions (categories) that contribute to outliers in measures (metrics) between two points in time. The key driver that contributes to an outlier helps you to answer the question: What happened to cause this anomaly?

If you are already using anomaly detection without contribution analysis, you can enable the existing ML insight to find key drivers. Use the following procedure to add contribution analysis and identify the key drivers behind outliers. Your insight for anomaly detection needs to include a time field and at least one aggregated metric (SUM, AVERAGE, or COUNT). You can include multiple categories (dimension fields) if you wish, but you can also run contribution analysis without specifying any category or dimension field.

You can also use this procedure to change or remove fields as key drivers in your anomaly detection.

To add contribution analysis to identify key drivers

1. Open your analysis and locate an existing ML insight for anomaly detection. Select the insight widget to highlight it.
2. Choose **Menu Options (...)** from the menu on the visual.
3. Choose **Configure anomaly** to edit the settings.
4. The **Contribution analysis (optional)** setting allows Amazon Quick Sight to analyze the key drivers when an outlier (anomaly) is detected. For example, Amazon Quick Sight can show you the top customers that contributed to a spike in sales in the US for home improvement products. You can add up to four dimensions from your dataset, including dimensions that you didn't add to the field wells of this insight widget.

To view a list of dimensions available for contribution analysis, choose **Select fields**.


If you want to change the fields you're using as key drivers, change the fields that are enabled in this list. If you disable all of them, Quick Sight won't perform any contribution analysis in this insight.

5. To save your changes, scroll to the bottom of the configuration options, and choose **Save**. To exit without saving, choose **Cancel**. To completely remove these settings, choose **Delete**.

Exploring outliers and key drivers with ML-powered anomaly detection and contribution analysis

You can interactively explore the anomalies (also known as outliers) in your analysis, along with the contributors (key drivers). The analysis is available for you to explore after the ML-powered anomaly detection runs. The changes you make in this screen aren't saved when you go back to the analysis.

To begin, choose **Explore anomalies** in the insight. The following screenshot shows the anomalies screen as it appears when you first open it. In this example, contributors analysis is set up and shows two key drivers.



The sections of the screen include the following, from top left to bottom right:

- **Contributors** displays key drivers. To see this section, you need to have contributors set up in your anomaly configuration.
- **Controls** contains settings for anomaly exploration.
- **Number of anomalies** displays outliers detected over time. You can hide or show this chart section.
- **Your field names** for category or dimension fields act as titles for charts that show anomalies for each category or dimension.

The following sections provide detailed information for each aspect of exploring anomalies.

Topics

- [Exploring contributors \(key drivers\)](#)
- [Setting controls for anomaly detection](#)
- [Showing and hiding anomalies by date](#)

- [Exploring anomalies per category or dimension](#)

Exploring contributors (key drivers)

If your anomaly insight is set up to detect key drivers, Quick Sight runs the contribution analysis to determine which categories (dimensions) are influencing the outliers. The **Contributors** section appears on the left.

Contributors contains the following sections:

- **Narrative** – At top left, a summary describes any changes in the metrics.
- **Top contributors configuration** – Choose **Configure** to change the contributors and the date range to use in this section.
- **Sort by** – Sets the sort applied to the results that appear below. You can choose from the following:
 - **Absolute difference**
 - **Contribution percentage** (default)
 - **Deviation from expected**
 - **Percentage difference**
- **Top contributor results** – Displays the results of the top contributor analysis for the point in time selected on the timeline at right.

Contribution analysis identifies up to four of the top contributing factors or key drivers of an anomaly. For example, Amazon Quick Sight can show you the top customers that contributed to a spike in sales in the US for health products. This panel appears only if you choose to include fields in contribution analysis when you configure the anomaly.

If you don't see this panel and you want to display it, you can turn it on. To do so, go to the analysis, choose anomaly configuration from the insight's menu, and choose up to four fields to analyze for contributions. If you make changes in the sheet controls that exclude the contributing drivers, the **Contributions** panel closes.

Setting controls for anomaly detection

You can find the settings for anomaly detection in the **Controls** section of the screen. You can open and close this section by clicking the word **Controls**.

The settings include the following:

- **Controls** – The current settings appear at the top of the workspace. You can expand this section by choosing the double arrow icon on the right side. The following settings are available for exploring outliers generated by ML-powered anomaly detection:
 - **Severity** – Sets how sensitive your detector is to detected anomalies (outliers). You should expect to see more anomalies with the threshold set to **Low and above**, and fewer anomalies when the threshold is set to **High and above**. This sensitivity is determined based on standard deviations of the anomaly score generated by the RCF algorithm. The default is **Medium and above**.
 - **Direction** – The direction on the x-axis or y-axis that you want to identify as anomalous. The default is [ALL]. You can choose the following:
 - Set to **Higher than expected** to identify higher values as anomalies.
 - Set to **Lower than expected** to identify lower values as anomalies.
 - Set to **[ALL]** to identify all anomalous values, both high and low.
 - **Minimum Delta - absolute value** – Enter a custom value to use to as the absolute threshold to identify anomalies. Any amount higher than this value counts as an anomaly.
 - **Minimum Delta - percentage** – Enter a custom value to use to as the percentage threshold to identify anomalies. Any amount higher than this value counts as an anomaly.
 - **Sort by** – Choose the method that you want to apply to sorting anomalies. These are listed in preferred order on the screen. View the following list for a description of each method.
 - **Weighted anomaly score** – The anomaly score multiplied by the log of the absolute value of the difference between the actual value and the expected value. This score is always a positive number.
 - **Anomaly score** – The actual anomaly score assigned to this data point.
 - **Weighted difference from expected value** – (Default) The anomaly score multiplied by the difference between the actual value and the expected value.
 - **Difference from expected value** – The actual difference between the actual value and the expected value (actual–expected).
 - **Actual value** – The actual value with no formula applied.
 - **Categories** – One or more settings can appear at the end of the other settings. There is one for each category field that you added to the category field well. You can use category settings to limit the data that displays in the screen.

Showing and hiding anomalies by date


The **Number of anomalies** chart shows outliers detected over time. If you don't see this chart, you can display it by choosing **SHOW ANOMALIES BY DATE**.

This chart shows anomalies (outliers) for the most recent data point in the time series. When expanded, it displays the following components:

- **Anomalies** – The middle of the screen displays the anomalies for the most recent data point in the time series. One or more graphs appear with a chart showing variations in a metric over time. To use this graph, select a point along the timeline. The currently selected point in time is highlighted in the graph, and includes a menu offering you the option to analyze contributions to the current metric. You can also drag the cursor over the timeline without choosing a specific point to display the metric value for that point in time.
- **Anomalies by date** – If you choose **SHOW ANOMALIES BY DATE**, another graph appears that shows how many significant anomalies there were for each time point. You can see details in this chart on each bar's context menu.
- **Timeline adjustment** – Each graph has a timeline adjustor tool below the dates, which you can use to compress, expand, or choose a period of time to view.

Exploring anomalies per category or dimension

The main section of the **Explore anomalies** screen is locked to the lower right of the screen. It remains here no matter how many other sections of the screen are open. If multiple anomalies exist, you can scroll out to highlight them. The chart displays anomalies in color ranges and shows where they occur over a period of time.



Each category or dimension has a separate chart that uses the field name as the chart title. Each chart contains the following components:

- **Configure alerts** – If you are exploring anomalies from a dashboard, select this button to subscribe to alerts and contribution analysis (if configured). You can set up the alerts for the level

of severity (medium, high, and so on). You can get the top five alerts for **Higher than expected**, **Lower than expected**, or **ALL**. Dashboard readers can configure alerts for themselves. If you open the **Explore Anomalies** page doesn't display this button if you opened the page from an analysis.

 Note

The ability to configure alerts is available only in published dashboards.

- **Status** – Under the **Anomalies** header, the status label displays information on the last run. For example, you might see "Anomalies for Revenue on November 17, 2018." This label tells you how many metrics were processed and how long ago. You can choose the link to learn more about the details, such as how many metrics were ignored.

Forecasting and creating what-if scenarios with Amazon Quick Sight

Using ML-powered forecasting, you can forecast your key business metrics with point-and-click simplicity. No machine learning expertise is required. The built-in ML algorithm in Amazon Quick Sight is designed to handle complex real-world scenarios. Amazon Quick Sight uses machine learning to help provide more reliable forecasts than available by traditional means.

For example, suppose that you are a business manager. Suppose that you want to forecast sales to see if you are going to meet your goal by the end of the year. Or, suppose that you expect a large deal to come through in two weeks and you want to know how it's going to affect your overall forecast.

You can forecast your business revenue with multiple levels of seasonality (for example, sales with both weekly and quarterly trends). Amazon Quick Sight automatically excludes anomalies in the data (for example, a spike in sales due to price drop or promotion) from influencing the forecast. You also don't have to clean and reprep the data with missing values because Amazon Quick Sight automatically handles that. In addition, with ML-powered forecasting, you can perform interactive what-if analyses to determine the growth trajectory you need to meet business goals.

Using forecasts and what-if scenarios

You can add a forecasting widget to your existing analysis, and publish it as a dashboard. To analyze what-if scenarios, use an analysis, not a dashboard. With ML-powered forecasting, Amazon

Quick Sight enables you to forecast complex, real-world scenarios such as data with multiple seasonality. It automatically excludes outliers that it identifies and imputes missing values.

Use the following procedure to add a graphical forecast to your analysis, and explore what-if scenarios.


Although the following procedure is for graphical forecasting, you can also add a forecast as a narrative in an insight widget. To learn more, see [Creating autonarratives with Amazon Quick Sight](#).

ML-powered forecasting is not compatible with [small multiples](#). To ensure accurate display of data and forecasts, avoid using small multiples in your visualizations.

To add a graphical forecast to your analysis

1. Create a visual that uses a single date field and up to three metrics (measures).
2. On the menu in the upper-right corner of the visual, choose the **Menu options** icon (the three dots), and then choose **Add forecast**.

Quick Sight automatically analyzes the historical data using ML, and displays a graphical forecast for the next 14 periods. Forecast properties apply to all metrics in your visual. If you want individual forecasts for each metric, consider creating a separate visual for each metric and adding a forecast to each.



3. On the **Forecast properties** panel at left, customize one or more of the following settings:
 - **Forecast length** – Set **Periods forward** to forecast, or set **Periods backward** to look for patterns to base the forecast on.

- **Prediction interval** – Set the estimated range for the forecast. Doing this changes how wide the band of possibility is around the predicted line.
- **Seasonality** – Set the number of time periods involved in the predictable seasonal pattern of data. The range is 1–180, and the default setting is **Automatic**.
- **Forecast boundaries** – Set a minimum and/or maximum forecast value to prevent forecast values from going above or below a specified value. For example, if your forecasting predicts the number of new hires the company will make in the next month to be in the negative numbers, you can set a forecast boundary minimum to zero. This stops the forecasted values from ever going below zero.

To save your changes, choose **Apply**.

If your forecast contains multiple metrics, you can isolate one of the forecasts by selecting anywhere inside the orange band. When you do this, the other forecasts disappear. Select the isolated forecast band again to have them reappear.

4. Analyze what-if scenarios by choosing a forecasted data point (in the orange band) on the chart, and then choosing **What-if analysis** from the context menu.

The **What-if analysis** panel opens at left. Set the following options:

- **Scenario** – Set a target for a date, or set a target for a time range.
- **Dates** – If you are setting a target for a specific date, enter that date here. If you are using a time range, set the start and end dates.
- **Target** – Set a target value for the metric.

Amazon Quick Sight adjusts the forecast to meet the target.

 Note

The **What-if analysis** option isn't available for multiple-metric forecasts. If you want to perform a what-if scenario on your forecast, your visual should contain only one metric.

5. Keep your changes by choosing **Apply**. To discard them, close the **What-if analysis** panel.

If you keep your changes, you see the new forecast adjusted for the target, alongside the original forecast without the what-if.

The what-if analysis is represented on the visual as a dot on the metric line. You can hover over the data points on the forecasting line to see the details.

Here are other things you can do:

- To interact with or remove a what-if analysis, choose the dot on the metric line.
- To create additional what-if scenarios, close the what-if analysis before choosing a new point on the line.

Note

What-if analyses can exist inside an analysis only, not inside a dashboard.

Generative BI with Quick Sight

Note

Powered by Amazon Bedrock: Amazon Q in Quick is built on Amazon Bedrock and includes [automated abuse detection](#) implemented in Amazon Bedrock to enforce safety, security, and the responsible use of AI.

Important

For the latest Amazon Quick chat experience — including cross-dataset Q&A with multi-dataset Topics, agentic workflows, and LLM-powered natural language querying — see [Ask questions, explore data, and get insights with chat in Amazon Quick](#). To learn about the new multi-dataset Topics that serve as the semantic layer for both chat and analysis, see [Working with Amazon Quick Sight Topics](#).

The Generative BI features described on this page represent the earlier experience. They continue to work as before.

With Amazon Quick chat, you can utilize the Generative BI authoring experience, create executive summaries of your data, ask and answer questions of data, and generate data stories.

To access all Quick Sight Generative BI features that are relevant to your task, choose the sparkle icon at the top right of any Quick page. In the pane that opens, the chat displays all content that is available based on the context of the task that you are performing. For example, if you're working in an Analysis, you can build a calculation, edit visuals, set up Q&A, or ask questions about your data. If you're working in a Dashboard, you can build a data story, generate an executive summary, or ask questions about the dashboard.

 Note

Generative BI features are not available in all AWS regions. To see a list of regions that Generative BI features are available in, see [Supported AWS Regions for Amazon Q in Quick](#)

Use the following topics to learn more about Generative BI.

Topics

- [Get started with Generative BI](#)
- [Augmenting Amazon Quick Sight insights with Amazon Q Business](#)
- [The Generative BI authoring experience](#)
- [Creating executive summaries](#)
- [Authoring Q&A](#)
- [Manage topic permissions through dashboards in Amazon Quick Sight](#)
- [Turn on the Dashboard Q&A experience in Amazon Quick Sight](#)
- [Q&A null support](#)
- [Improve Q&A accuracy with custom instructions](#)
- [Asking and answering questions of data with Generative BI](#)
- [Opting out of Generative BI](#)
- [Working with Amazon Quick Sight Topics](#)
- [Working with data stories in Amazon Quick Sight](#)
- [Working with scenarios in Amazon Quick Sight](#)

Get started with Generative BI

To get started with Quick Sight Generative BI capabilities, upgrade your account's users to Admin Pro, Author Pro, or Reader Pro roles. Pro roles grant users access to all Generative BI capabilities that are relevant to the role that's assigned to the user. Pro users can share generative Q&A legacy Topics with another user. To understand which Generative BI capabilities are available to the different user roles in Quick, see the table below. To understand how subscription names map to user roles, see [Understanding Amazon Quick subscriptions and roles](#).

 **Note**

Non-Pro Authors and Readers can still access Generative Q&A legacy Topics if an Author Pro or Admin Pro user shares the legacy Topic with them. Non-Pro Authors and Readers can also access data stories if a Reader Pro, Author Pro, or Admin Pro shares one with them.

Feature name	Feature description	Reader	Author	Admin	Reader Pro	Author Pro	Admin Pro
Creating a data story with Generative BI	Build data stories that explain your data with visuals, insights, and ideas to help improve your business.	⊗ No	⊗ No	⊗ No	✔ Yes	✔ Yes	✔ Yes
Viewing a	View narrative	✔ Yes*	✔ Yes*	✔ Yes*	✔ Yes	✔ Yes	✔ Yes

Feature name	Feature description	Reader	Author	Admin	Reader Pro	Author Pro	Admin Pro
generated data story in Amazon Quick Sight	data stories that are shared with you.						
Authoring Q&A	Create and refine legacy Topics that utilize Generative Q&A for QuickSight dashboards.	⊗ No	⊗ No	⊗ No	⊗ No	✓ Yes	✓ Yes
Asking and answering questions of data with Generative BI	Ask questions about data to accelerate data driven decisions with multi-visual answers.	✓ Yes*	✓ Yes*	✓ Yes*	✓ Yes	✓ Yes	✓ Yes

Feature name	Feature description	Reader	Author	Admin	Reader Pro	Author Pro	Admin Pro
Creating executive summaries	Get an executive summary of key insights from a Quick Sight dashboard.	⊗ No	⊗ No	⊗ No	✓ Yes	✓ Yes	✓ Yes
The Generative BI authoring experience	Create an analysis to build visuals, calculations, and refine existing visuals with natural language.	⊗ No	⊗ No	⊗ No	⊗ No	✓ Yes	✓ Yes

*Non-pro roles in accounts that were created on or after April 30, 2024 can access Q&A legacy Topics that are shared with them. If your Quick account was created before April 30, 2024 and you want to opt-in to this new feature, contact your AWS account team.

Any Quick administrator can upgrade a user to a Pro role with the following procedure.

To upgrade a user to a Pro role

1. Open the [Quick console](#).

2. Choose the user icon at the top right, and then choose **Manage Quick**.
3. Choose **Manage users** to open the **Manage Users** page.
4. To change the role of an existing user, locate that user on the **Manage Users** table and choose the role that you want to grant them from the **Role** dropdown.

For more information about managing Quick users, see [Managing user access inside Amazon Quick](#).

Augmenting Amazon Quick Sight insights with Amazon Q Business

Amazon Quick account admins can connect their Quick account to Amazon Q Business to augment insights with unstructured data sources. [Amazon Q Business](#) is a generative AI assistant that helps your team work smarter. It can answer questions, provide summaries, generate content, and securely complete tasks based on the information in your enterprise systems.

When an Quick account is integrated with Amazon Q Business, users can now leverage this vast repository of organizational knowledge alongside their structured data analytics. This integration allows for more comprehensive and context-rich insights, as it combines quantitative data from Quick with qualitative information from various business documents and applications.

For more information about connecting your Amazon Q Business account with Quick, see [Creating an Quick-integrated application](#).

Use the following topics to configure an Amazon Q Business application in Quick.

Topics

- [Considerations](#)
- [Configuring an Amazon Q Business application in Amazon Quick Sight](#)
- [Connect a Quick account to an existing Amazon Q Business application](#)
- [Disconnect an Amazon Q Business application from an Amazon Quick account](#)

Considerations

The following limitations apply to the Amazon Q Business application.

- Quick and Amazon Q Business must exist in the same AWS account. Cross account calls are not supported.

- Quick and Amazon Q Business accounts need to exist in the same AWS Region. Cross Region calls are not supported. For a list of all supported Quick Regions, see [Supported AWS Regions for Amazon Q in Quick](#). For a list of all supported Amazon Q Business Regions, see [Service quotas for Amazon Q Business](#).

If your Quick account exists in more than one Region, you can connect one Amazon Q Business application from each Region to the Quick account. For example, if your Quick account exists in US East (N. Virginia) and US West (Oregon), one Amazon Q Business application located in US East (N. Virginia) and one Amazon Q Business application located in US West (Oregon) can be connected to the Quick account.

- Quick and Amazon Q Business accounts that are integrated need to use the same identity methods. For example, if a Quick account uses IAM Identity Center for identity management, the Amazon Q Business account that it is integrating with must also use IAM Identity Center for identity management.
- Email addresses that are associated with Quick users and groups are used to perform authorization checks in Amazon Q Business.

Configuring an Amazon Q Business application in Amazon Quick Sight

Use the following procedure to connect an Amazon Quick account with Amazon Q Business

1. Open the [Quick console](#).
2. Choose the user icon at the top right, and then choose **Manage Quick**.
3. Choose **Security & permissions**.
4. On the **Quick access to AWS services** page, choose the **Amazon Q Business application** checkbox.
5. On the **Create an Amazon Q Business connection to unstructured data** popup that appears, choose the Quick Region that you want your connection to be in.
6. Choose **Done**.
7. When you choose **Done**, your Amazon Q Business account is created and you are redirected to a new tab that shows the **Applications** page of the Amazon Q Business console.
8. For **Applications**, choose the Amazon Q Business connection that you created in Quick.
9. The **Application details** page of your connection opens. Choose the **Index** tab, and then choose **Select index**.

10. In the popup that appears, choose the **Index provisioning** option that you want to use, and then choose **Confirm**. For more information about indexes in Amazon Q Business, see [Creating a retriever for an Amazon Q Business application](#).
11. After you choose an index, set up a data source connection. To set up a data source connection, choose the **Data sources** section of the **Enhancements** menu in the left side pane.
12. Choose **Add data source**.
13. Choose the data source that you want to add. The data source that you choose determines the steps that are required to configure the data source connection. For more information about adding a data source to an Amazon Q Business account, see [Connecting Amazon Q Business data sources](#). When you finish setting up the data source configuration, choose **Add data source**.

After you choose an index, a retriever, and a data source for your Amazon Q Business account, your connection to Amazon Q Business is complete and you can return to the Quick console.

Connect a Quick account to an existing Amazon Q Business application

If you already have an Amazon Q Business application that uses the same identity management and exists in the same Region as your Quick account, use the following procedure to link the existing Amazon Q Business account to Quick.

1. Open the [Quick console](#).
2. Choose the user icon at the top right, and then choose **Manage Quick**.
3. Choose **Security & permissions**.
4. On the **Quick access to AWS services** page, choose the **Amazon Q Business application** checkbox.
5. On the **Create an Amazon Q Business connection to unstructured data** popup that appears, choose the Quick Region that you want your connection to be in.
6. Choose your existing Amazon Q Business application from the dropdown.

Note

Your Amazon Q Business application does not appear if the application exists in a different Region than your Quick account or if the application uses a different identity management option than your Quick account.

After you choose your Amazon Q Business application from the dropdown, the connection between Quick and Amazon Q Business is configured.

Disconnect an Amazon Q Business application from an Amazon Quick account

Quick account admins can use the following procedure to disconnect an Amazon Q Business application from a Quick account.

1. Open the [Quick console](#).
2. Choose the user icon at the top right, and then choose **Manage Quick**.
3. Choose **Security & permissions**.
4. On the **Quick access to AWS services** page, choose **SELECT APPLICATION**.
5. Perform one of the following options:
 - a. To disconnect a single Amazon Q Business application from a Quick account, navigate to the application that you want to remove, open the dropdown, and choose **NONE**.
 - b. To disconnect all Amazon Q Business applications from a Quick account, uncheck the **Amazon Q Business application** checkbox.

When you disconnect an Amazon Q Business application from a Quick account, the Amazon Q Business application that you created for Quick is not deleted. The application, index, retriever, and any unstructured data source connections that you configured remain in your Amazon Q Business account.

The Generative BI authoring experience

With Quick chat, authors can use new Generative BI capabilities to build calculated fields and to build and refine visuals. You can also generate complete multi-sheet analyses from natural language prompts. For more information, see [Generating an analysis with natural language prompts](#).

Use the following topics to learn more about the Generative BI authoring experience.

Topics

- [Build visuals with Generative BI](#)
- [Build calculations with Generative BI](#)
- [Refine visuals with generative BI](#)

Build visuals with Generative BI

Quick authors can use the **Build a visual** button to build a custom visual that's generated from author input. The author's input uses natural language to describe the desired outcome for the new visual. You can enter a custom description, or you can choose from a list of generated suggestions that Amazon Q has generated for the topic that's attached to the analysis. The following image shows a custom visual that's created with the **Build a visual** menu.

To build a visual with Generative BI

1. Navigate to the analysis that you want to work in and choose **Ask to build a visual**.
2. In the **Build a visual** panel that appears, perform the following steps.
 - a. Describe the data that you want to visualize. You can enter a custom description, or you can choose from the **Suggested** questions that are generated based on the analysis' data.


When you describe the data that you want to visualize, you can phrase it as a question, or you can use conversational phrases or filters. For example, you can enter "How many people signed up for a free trial last month?" or "Free trial sign ups by month." Both statements generate a visual that shows the number of free trial sign-ups by month. You can also get responses to vague language or keyword style requests.

Suggested questions can include a mix of artificial intelligence (AI) generated questions and human verified questions. Human verified questions appear with a check mark next to the suggestion.

- b. Choose **Build**.
- c. Review the visual that generates. To refine the data presented in the visual, enter a new description into the **Build** bar, and then choose **Build**. Use the forward and back arrows to review the changes made to the visual without losing any progress.
- d. When you're satisfied with the visual, choose **ADD TO ANALYSIS**.

Build calculations with Generative BI

With Generative BI, you can use natural language prompts to create calculated fields in Amazon Quick Sight, as shown in the following image. For more information about calculated fields in analyses, see [Adding calculated fields](#).



To build a calculated field with Generative BI

1. Navigate to the analysis that you want to work in and choose **Data** from the toolbar at the top of the page. Then choose **Add calculated field**.
2. In the calculation editor that appears, choose **Build**.
3. Describe the calculation outcome that you want to achieve. For example, "year over year percent change in daily sales."
4. Choose **BUILD**.
5. Review the expression that's returned, and then choose **Insert** to add it to the expression editor. You can also choose the **Copy** icon to copy the expression to your clipboard. To delete the expression and start over, choose the **Delete** icon next to the expression.
6. When you're finished, close the editor.

After you add a calculation to the expression editor, you must name the calculation before you can save it.

Refine visuals with generative BI

Quick authors can also use natural language prompts to edit visuals in an analysis, as shown in the following visual. Authors can use this functionality to edit visuals without performing manual tasks in the Quick UI. Authors can only use Generative BI to perform formatting tasks that are currently supported in Quick.

The following types of edit are supported:

- Change a visual's type.
- Show or hide axis titles, axis labels, or data labels.
- Show, hide, or change the title of a chart.
- Change axis and table column names.
- Add fields or field wells to a visual.

- Remove fields from a visual.
- Change the aggregation of an axis.
- Show or hide legends and grid lines.
- Show or hide data zoom.
- Add fields or field wells to a visual.
- Change or remove a visual's sort controls.
- Update the conditional formatting of a visual's colors, color gradients, background color, or text color.
- Change the time granularity of a visual.
- Adjust axis scaling and range, as well as maximum and minimum values.
- Change font sizes of titles and subtitles.
- Show, hide, and adjust data labels.
- Adjust column formatting (change between number, percent, date, and currency).

To edit a visual with Generative BI

1. Navigate to the visual that you want to edit, and then choose **Edit with Q**.
2. Describe the task that you want performed, and then choose **APPLY**.
3. Review the visual changes. If you're satisfied with the generated changes, close the **Edit visual** modal. To undo the changes, choose **Undo** and enter a new prompt.

Creating executive summaries

With Quick chat, you can leverage large language models (LLMs) to generate executive summaries of dashboards. Executive summaries are based on Quick Sight's suggested insights for a dashboard. Executive summaries help readers find key insights at a glance without the need to pinpoint specific data from a dashboard's visuals.

To turn on executive summaries for a dashboard, turn on **Allow executive summary** on the **Publish a dashboard** modal.

For more information about how readers can interact with executive summaries, see [Generate an executive summary of an Amazon Quick Sight dashboard](#).

Executive summaries work best when an analysis has multiple suggested insights. To see a list of all suggested insights for an analysis, navigate to the analysis that you want to work in, and then open the **Insights** pane.

Authoring Q&A

Note

The topics referenced in this section refer to legacy Topics. For information about the new multi-dataset Topics experience, which supports cross-dataset relationships and runtime joins for both chat and analysis, see [Working with Amazon Quick Sight Topics](#).

Converting to the Generative Q&A experience

If you have existing topics, you can easily convert these to leverage our new generative capabilities. Navigate to a topic, and then choose **Convert** next to the topic name. You will then be prompted to **Duplicate & Convert Topic** in a dialog box. We duplicate your topic for you so that the conversion to our beta experience does not impact your end users. Once you are satisfied with topic performance in the new experience, you can unshare the original topic and share the new one.

Named entities

Named entities are one of the most important components of topic curation. The information contained in named entities — specifically, the ordering of fields and their ranking — is what makes it possible to present contextual, multi-visual answers in response to even vague questions. Authors can find named entities by navigating to a topic, choosing the **Data** tab, and then choosing the **Named Entities**. From here, authors can preview or edit existing named entities, and create new ones.

Authors can configure the following facets of named entities:

1. **Fields:** Choose a dataset, and then choose which fields from that dataset to include. This defines the scope of data that will be considered when using this named entity to answer enduser questions.
2. **Field Rank and Presentation:** The relative rank of the dimensions and measures in a named entity determines how those fields are used when generating contextual, multi-visual answers. Note in the following demo that adjusting the relative rank of **Profit** so that it is higher than

Sales leads to different data being displayed. By default, the order of fields in the table visual is the same as the field rank. However, you can control these two individually by turning off **Sync table view with field ranking**.

3. **Show / Hide in Presentation:** Fields that are included in named entities can simultaneously be hidden from the tabular presentation of the named entity, while still providing additional context in other components of the answer.

Measure aggregations

Authors have fine-grained control over aggregated measures in topics. Across Quick Sight, measures are defaulted to SUM, unless they have custom aggregations defined in a calculated expression. To change this, navigate to the measure in the list of data fields, and specify a different default aggregation. You can also disallow aggregations, which will prevent them from being applied even if a user specifically asks for them. Lastly, you can specify that a measure is non-additive. This is useful for pre-computed metrics, such as percentages, which should not be re-combined in any way. Doing so will force MEDIAN or AVG depending on your use case.

Manage topic permissions through dashboards in Amazon Quick Sight

Note

The topics referenced in this section refer to legacy Topics that are linked to dashboards for Q&A. For information about the new multi-dataset Topics experience, see [Working with Amazon Quick Sight Topics](#).

Quick enables Authors to manage permissions for dashboards and their linked topics from a single location. When sharing dashboards with Q&A enabled, Authors can control topic viewer access directly from a dashboard's sharing preferences, eliminating the need to manage permissions in multiple locations.

To enable Q&A on a dashboard with a linked topic:

1. Open the [Quick console](#).
2. Open the analysis of the dashboard with Q&A enabled and topic linked that you want to publish.
3. Choose **Publish**.

4. Check the **Allow data Q&A** check box.
5. Choose **MANAGE Q&A** and select **Use a linked topic for Build visual and Q&A**.
6. Select the desired linked topic from the dropdown menu.
7. Choose **APPLY CHANGES**, then choose **Publish dashboard**.

To conveniently manage topic access from a dashboard:

1. Open the [Quick console](#).
2. Open the dashboard with a linked topic that you are a co-owner of.
3. Select the share icon and choose **Share dashboard**.
4. In the row of your selected user, flip on/off the **Share as "topic viewer"** toggle to grant/revoke viewer access to the linked topic.
5. In the row of your selected shared folder, flip on/off the **Add topic to folder** toggle to add/remove the linked topic to/from the shared folder.

To share the dashboard and its linked topic to all users and groups:

1. Open the [Quick console](#).
2. Open the dashboard with a linked topic that you are a co-owner of.
3. Select the share icon and choose **Share dashboard**.
4. On the bottom-left of the panel, under **Auto-share linked topic for**, flip the **All dashboard users and groups** toggle on. This will grant viewer access to the linked topic when the dashboard is shared. Flip the toggle off to cancel this behavior.

After the dashboard with a linked topic has been shared, users will immediately be able to ask questions about their data. Navigate to **Ask a question about <topic name>** at the top of the dashboard to start asking questions.

Turn on the Dashboard Q&A experience in Amazon Quick Sight

Quick allows any Author to enable Q&A directly from their dashboards in one click without the need to create a Topic in Quick Sight. To do this, publish your dashboard and check the **Allow data Q&A** checkbox from the dashboard publishing menu. When you turn on dashboard Q&A, you can choose which datasets to use for dashboard Q&A to ensure that your end users get the answers they need.

Dashboard Q&A queries all rows and columns in the included datasets - beyond what is visible in the dashboard. To protect sensitive or confidential data, enable [row-level security \(RLS\)](#) and/or [column-level security \(CLS\)](#).

The following table compares feature availability between dashboard Q&A and topic Q&A.

Q&A feature	Dashboard Q&A	Topic Q&A
Allows users in all roles to ask and answer questions of data	✔ Yes	✔ Yes
Allows author and admin roles to enable data Q&A on dashboards	✔ Yes	✘ No (Pro users only)
Supported in Quick console embedding	✘ No	✔ Yes
Ability to add reviewed answers	✘ No	✔ Yes
Ability to customize Q&A-specific metadata	✘ No	✔ Yes
Ability to support autocomplete for data values	✘ No	✔ Yes

Use the procedure below to enable dashboard Q&A on a Quick Sight dashboard.

To enable dashboard Q&A on a dashboard

1. Open the [Quick console](#).
2. Open the analysis that the dashboard that you want to publish with Q&A enabled.
3. Choose **Publish**.
4. Check the **Allow data Q&A** check box.
5. (Optional) Choose **MANAGE Q&A** to choose which datasets you want to include in the dashboard Q&A experience. By default, all datasets that are used by the dashboard are included.

6. Choose **APPLY CHANGES**, and then choose **Publish dashboard**.

After you publish a dashboard with the dashboard Q&A experience enabled, users can ask questions about their data with the **Ask a question about this dashboard** input at the top of the dashboard.

Quick allows any user to ask questions on dashboards that have dashboard Q&A enabled. However, dashboard Q&A is a feature that incurs the associated enablement fee. Quick admins can disable this feature at the account level at any time. Use the following procedure to disable dashboard Q&A across an entire Quick account.

1. Open the [Quick console](#).
2. Choose the user icon in the top right, and then choose **Manage Quick**.
3. Choose **Security & permissions**.
4. Navigate to the Amazon Q section, and then choose **Manage**.
5. Toggle **Manage Dashboard Q&A** off.

When you toggle **Manage Dashboard Q&A** off, dashboard Q&A is removed from any dashboards that have dashboard Q&A enabled. If your Quick account does not have Pro users or topics, this action stops the Amazon Q enablement fee from billing your Quick account. This setting does not impact Pro users or existing topics in Quick. For more information about opting out of Generative BI, see [Opting out of Generative BI](#).

Q&A null support

Amazon Quick Sight Q&A has comprehensive support for null value handling, enabling users to create more sophisticated analyses and answer complex business questions. This functionality allows for precise filtering of null values, intuitive queries about missing data, and dynamic chart interactions.

Add a filter to include or exclude null values

To add a filter to include or exclude null values

1. Open the [Quick console](#).
2. Choose **Topics** and then open the topic you want to add a filter to.

3. Choose the **Data** tab.
4. Under **Data Fields**, choose **Add filter**.
5. On the **Filter configuration** page that opens, do the following:
 - a. For **Name**, enter a name for the filter.
 - b. For **Dataset**, choose a dataset that you want to apply the filter to.
 - c. For **Field**, choose the field that you want to filter for.
 - d. For **Null Option**, choose one of the dropdown options:
 - **No null option selected** - No option is selected to filter nulls.
 - **Include nulls only** - Filter for only nulls on the field selected.
 - **Exclude nulls only** - Filter for only non nulls on the field selected.
 - e. (Optional) To specify when the filter is applied, choose **Apply the filter anytime the dataset is used**, and then choose one of the following:
 - i. **Apply always** - Filter is applied whenever a column from the specified dataset is linked to a question.
 - ii. **Apply always, unless a question results in an explicit filter from the dataset** - Filter is applied whenever a column from the specified dataset is linked to a question, unless the question contains its own explicit filter for the same field.
 - f. Choose **Save**.

The filter is added to the list of fields in the topic. You can edit the description for it or adjust it when the filter is applied.

Ask a question on null values


You can use Q&A to directly ask questions about null values, such as:

- What is the total sales amount for records where the segment is null?
- Display accounts without assigned representatives.
- List projects with no completion date.
- Show inventory items without category assignments.
- What percentage of total orders have non null values in the license field by segment?
- Which orders do not have a customer assigned?

Manage null values in visualizations

After generating visualizations through the Q&A bar, you can interact with the charts using various null value actions, including focusing only on null values or excluding null values. These chart actions help you analyze and filter your data dynamically based on null value presence.


Choose either **Focus only on null** or **Exclude null** to appropriately filter the results.



Refine query interpretations for null value handling

Once the visualizations are generated based on your query, you can adjust how null values are handled.

1. Locate the **Interpreted as** section below your query.
2. Select the field you wish to modify.
3. From the dropdown menu, choose **Null Options** to adjust null value handling.



For categorical fields, empty values are not the same as null values. To convert empty values into nulls:

1. Open the [Quick console](#).
2. Choose **Topics** and then open the topic you want to add a filter to.
3. Choose the **Data** tab.
4. Choose **Add calculated field**.
5. Enter a name for the **Add name** field.
6. Choose a categorical field and enter an expression to convert empty values to null values: `ifelse({Segment}="", NULL, {Segment})`.
7. Choose **Save**.

Improve Q&A accuracy with custom instructions

Custom Instructions enables Authors to curate Amazon Q's responses to questions by adding domain-specific knowledge that can't be captured through a topic's metadata settings, such as synonyms or semantic types. By providing these metadata descriptions or custom instructions, Authors can guide Amazon Q to align its responses with distinct definitions, preferences, and expert knowledge—ensuring more accurate, relevant, and tailored answers that are better suited for their business needs.

Use the following table to understand when and how to apply different types of metadata to improve Q&A answer accuracy. Each metadata type plays a unique role in clarifying context,

resolving ambiguity, and ensuring that answers are aligned with business rules or domain-specific terminology.

Metadata Type	When to Use	How it Improves Answer Accuracy
Field-Level Description	When the Q&A system needs to understand ambiguous or domain-specific column names (for example, DTC Spend).	Clarifies field semantics so the model can answer more precisely (for example, interpreting DTC Spend as Direct-to-Consumer marketing expense).
Topic-Level Description	When users may ask broad or ambiguous questions and Amazon Q needs more context about the topic's overall purpose (for example, sales performance vs. clinical trial data).	Helps disambiguate general terms and steer answers toward the right domain (for example, sales vs. marketing).
Dataset Description	When users have access to multiple datasets and the Q&A system needs to identify which one best fits the question.	Enables dataset selection logic by providing context about each dataset's purpose and content.
Topic-Level Custom Instructions	When a topic has specific business rules, timeframes, or definitions (for example, fiscal year ≠ calendar year).	Applies custom logic or definitions (for example, defining Q1 as August-October) to tailor answers appropriately.

Adding field-level descriptions

To add field-level descriptions:

1. Open the [Quick console](#).
2. Choose **Topics** and then open the topic you want to add descriptions for.
3. From the topic details page, select the **Data** tab then choose the **Data Fields** sub-tab.
4. Add a description to improve answer accuracy for each included field. This is especially important for field names that contain bespoke corporate knowledge to understand.

If you have multiple date fields, for example, clear descriptions can help Amazon Q distinguish between them and choose the most relevant one based on the user’s question. In the sample below, an Author added descriptions for **Solution Create** and **Topic Create**, which enables Amazon Q to more accurately select the appropriate date field in context.

Include

☐

Friendly name



Solution Create



Field name: solution_create

Field type: Time

Need help understanding these settings? [Learn more](#)

Description

The timestamp when an approver user marks their reply as a solution on a post.

Include	Friendly name
☐ 	 Topic Create 

Field name: topic_create Field type: Time Need help understanding these settings? [Learn more](#)

Description

The timestamp when a new Community tag is first created. Tags categorize posts into meaningful buckets for the product team, like AI/ML, calculated fields, data connectors, etc.

Adding topic-level descriptions

To add topic-level descriptions:

1. Open the [Quick console](#).
2. Choose **Topics** and then open the topic you want to add descriptions for.
3. From the topic details page, select the **Summary** tab.
4. Under **Topic Details**, add a description to provide more context about the topic's overall purpose.

Adding dataset descriptions

To add dataset descriptions:

1. Open the [Quick console](#).
2. Choose **Topics** and then open the topic you want to add descriptions for.
3. From the topic details page, select the **Data** tab then choose the **Datasets** sub-tab.
4. Add a description to help improve dataset selection logic.

Adding topic-level custom instructions

To add custom instructions:

1. Open the [Quick console](#).
2. Choose **Topics** and then open the topic you want to add descriptions for.
3. From the topics details page, select the **Custom Instructions** tab.
4. Add topic-level guidance to help the chat better understand the context, terminology, or intent that is specific to the selected topic. This can include disambiguation tips, field relationships, definitions for terms that can't be captured in a calculated field or topic filter, or instructions for customizing relative date ranges.

Best practices for writing custom instructions

Match cell values precisely

- Use the exact cell value from the database, including casing and formatting.
- If the value is ambiguous, reference its source column to clarify.

Examples:

- Instead of: *"AMZ are Amazon customers"*

Use: *"AMZ are 'Amazon.com, Inc.' customers"*

- Instead of: *"ETPs are enterprise customers"*

Use: *"ETPs are customers from the enterprise Segment"*

Be specific and quantitative

Avoid vague language—be clear about filters, thresholds, and source columns.

Example:

- Instead of: *"Filter large customers when talking about sales"*

Use: *"Filter customers where Annual Revenue > \$1M when talking about sales"*

Use formatting for clarity, not function

Spacing and line breaks do not affect model behavior, but help authors read and maintain instructions more easily.

Understand what custom instructions cannot do

Custom instructions improve the understanding of your business context, but they do not add new capabilities. These instructions will not:

- Change chart type selections
- Perform calculations or fill nulls
- Create new fields
- Control formatting, colors, or legends
- Alter the narrative or number/type of visuals

Adding field-level descriptions in data preparation for dashboard-based Q&A

In addition to Topic-based descriptions, you can create field-level definitions to enhance [Dashboard Q&A](#) functionality. Adding specific definitions to individual fields during the data preparation phase improves the answer accuracy when users ask questions about particular dashboard elements.

To add field-level descriptions for dashboard-based Q&A:

1. Open the [Quick console](#).
2. Choose **Data**, open a dataset that you have access to, and select **EDIT DATASET**.
3. For each relevant field, choose the three-dot menu and select **Edit name & description**.
4. Add a description to enhance answers for dashboard-related questions.
5. Choose **Apply** to save your changes.

Asking and answering questions of data with Generative BI

Note

The Q&A experience described in this section uses legacy Topics. To view the multi-visual experience, the legacy Topic author must add named entities and convert an existing topic

to use generative capabilities or create a new generative topic. For more information, see [Authoring Q&A](#).

For the new multi-dataset Topics experience, which uses an LLM-powered chat agent to generate cross-dataset SQL queries with runtime joins, see [Using Topics in Amazon Quick chat](#).

Accelerate data-driven decisions with humanistic Q&A that includes:

- AI-generated narrative that highlights key insights
- Multi-visual answer that provides the answer to your question along with supporting visuals to add valuable context
- Home page for every topic with AI-generated and author-reviewed suggested questions and automated data previews to see what data you can ask about

Choose the sparkle icon at the top right. Once you open your topic, there is a home page with a list of suggested questions and **What's in your topic** to see what data you can ask about.

When there are multiple dates available, choose **more...** to view them. For example, in this Student Enrollment Trends topic, there is data available for enrollment data spanning from 2018 to 2023, but there is also student Date of Birth (DOB) data ranging from 1973 to 2005.

Choose a suggested question or type your own question to get started. By hovering over a sentence in the AI-generated narrative, you can clearly identify the source visualization and verify the values. Each visualization is interactive and can be added to your pinboard.

You can get answers to a variety of questions from vague to precise.

If you don't have a precise question in mind, you can ask a vague question that is only one word or a short phrase, like *"sales"* or *"top students."* You can include additional filter criteria with these vague questions like *"top students last semester."*

Question examples include:

- Entity name: *"Order Details"*

- **Note**
You can find the entities from the topic home page and in the **What's in *topic*** tab at the top of the list.
- Field name: "Segment"
- Field values: "Acme Inc.," "Washington DC"
- Vague (or implicit) filters: "best account managers," "bottom products"

For precise questions that are supported, see this table of question types: [Types of questions supported by Q](#). Examples include "product with largest WoW growth %" or "forecast sales for APAC customers by quarter." It covers a range of filters, like top/bottom, relative and absolute date filters, period-to-date and period-over-period, and more. It also supports analytical questions, like percent of total, or "why did sales drop in October 2023?"

- **Tip**
To help you form questions, think *Who, What, Where, When* and *Why*.


Unpacking your answer:

- **Interpreted as:** – This is how Amazon Q interpreted your question. It will map your words to the underlying data so you can verify that you were correctly understood. If not, adjust your question or leave feedback for your author.
- **AI-generated narrative:** – A summary of the visuals that highlights key insights. If your Quick account is connected to an Amazon Q application, you may receive additional insights from unstructured data sources under **Insights from Q Business**. You can see the unstructured sources that are used in the **Sources** collapsible. For more information about connecting a Quick account to a Amazon Q Business application, see [Augmenting Amazon Quick Sight insights with Amazon Q Business](#).
- **Visuals:** – Visuals consist of: center visual that directly answers the question, supporting visual on the right that provides context, relevant KPIs, and a details table at the bottom.

Note

If the field is not included in a named entity, then it will display as a single visual.

- **Did you mean:** – When there are multiple interpretations to your question, it will display a list of alternate answers that you can select to align with your intended question.
- In the following example, the question "top customers" can be interpreted in several ways, including by "Total Sales," "Total Profit," or "number of customers."



Other tips

- To re-size the panel, drag the left side.
- Add important visuals to your pinboard for quick access. View your pinboard from the top of the Amazon Q pane.
- Provide feedback for your topic author to see and make improvements.

Opting out of Generative BI

Quick accounts are charged if Generative BI is active in the account. Generative BI is considered active if your account uses any of the following capabilities:

- Pro users
- Topics

- Dashboard and visual indexing
- Dashboard Q&A

To avoid being charged for Generative BI by completely deactivating it, perform the following steps.

Warning

Opting out of Generative BI will disable AI-powered features and stop related charges. This process involves:

- Removing or changing Pro user roles to standard roles
- Deleting all topics in your account
- Disabling dashboard indexing and Q&A features

Before proceeding: Review the steps carefully and ensure you understand which features will be disabled.

To opt out of Generative BI

1. Ensure there are no Pro users or user groups mapped to Pro roles in the account by performing the following steps:
 - To update or remove Pro users using APIs:
 - If you use Quick identity (with or without IAM federation):
 - a. Find users that have Pro roles using the [ListUsers](#) API.
 - b. Either change the users' roles using the [UpdateUser](#) API, or remove the users from the account using the [DeleteUser](#) API.
 - If you use IAM Identity Center or Microsoft Active Directory:
 - a. Find group of users mapped to Pro roles using the [ListRoleMemberships](#) API.
 - b. Create new user groups with the same users, but mapped to different roles, using the [CreateRoleMemberships](#) API.
 - c. Delete the previous user groups mapped to Pro roles using the [DeleteRoleMemberships](#) API.

- To update or remove Pro users using the Quick console:
 - a. Open the [Quick console](#).
 - b. Choose the profile icon, then choose **Manage Quick**.
 - c. If necessary, in the left navigation pane, choose **Manage users**.
 - If you use Quick identity (with or without IAM federation), update user roles or delete users using the steps in [Viewing Amazon Quick account details](#) or [Deleting a Amazon Quick user account](#).
 - If you use IAM Identity Center or Microsoft Active Directory, update group and role mappings or delete user groups using the steps in [Managing user access](#).
- 2. Ensure there are no topics in the account by performing the following steps:
 - a. Use the [ListTopics](#) API to list all topics in the account for each AWS Region where topics are used.
 - b. For each topic, do one of the following:
 - If you are an owner or co-owner of the topics, delete the topics using the [DeleteTopic](#) API.
 - If you're not an owner or co-owner of the topics:
 - Identify the owners of each topic using the [DescribeTopicPermissions](#) API, then ask them to delete their topics using the [DeleteTopic](#) API.
 - Make yourself a co-owner of the topics using the [UpdateTopicPermissions](#) API , then delete the topics using the [DeleteTopic](#) API.
- 3. Ensure that dashboard and visual indexing and Dashboard Q&A are disabled by performing the following steps:
 - To disable dashboard and visual indexing and Dashboard Q&A using APIs:
 - a. Disable dashboard and visual indexing using the [UpdateQuickSightQSearchConfiguration](#) API.
 - b. Disable Dashboard Q&A using the [UpdateDashboardsQAConfiguration](#) API.
 - To disable dashboard and visual indexing and Dashboard Q&A using the Quick console:
 - a. Open the [Quick console](#).
 - b. Choose the profile icon, then choose **Manage Quick**.
 - c. Under the **Account** section, choose **Amazon Q**.

- d. Disable each of the options.

Working with Amazon Quick Sight Topics

Applies to: Enterprise Edition

Intended audience: Amazon Quick administrators and authors

A *Topic* in Quick Sight is the multi-dataset semantic layer that brings multiple enriched datasets together into a unified data model. Topics let you define relationships between datasets so that Quick Sight can perform runtime joins across them — whether you're building analysis visuals or asking natural language questions through Amazon Quick chat.

With Topics, you can:

- Add up to 12 datasets to a single Topic and define relationships (join keys) between them.
- Add custom instructions that guide the AI engine on cross-dataset logic, disambiguation rules, and business definitions.
- Use the Topic as a data model in Quick Sight analysis sheets to build visuals with fields from multiple datasets — Quick Sight performs runtime inner joins automatically.
- Use the Topic in Amazon Quick chat where the LLM-powered chat agent traverses relationships across datasets, generates cross-dataset SQL with appropriate joins, and returns unified answers.
- Preserve each dataset at its native granularity — no pre-joining or denormalization required.

Each dataset in a Topic should be independently enriched with semantic metadata (column descriptions, synonyms, semantic types, and custom instructions) before being added to the Topic. For more information about enriching datasets, see [Data Preparation Experience \(New\)](#).

Note

If you have existing Topics that were created before this release, they are now classified as *legacy Topics*. Legacy Topics continue to work as before. For more information, see [Working with legacy Topics](#).

To learn more about working with Topics, continue to the following sections.

Topics

- [How Topics work](#)
- [Creating Quick Sight Topics](#)
- [Defining relationships between datasets in a Topic](#)
- [Adding custom instructions to a Topic](#)
- [Using Topics in Quick Sight analysis](#)
- [Using Topics in Amazon Quick chat](#)
- [Sharing Quick Sight Topics](#)
- [Working with legacy Topics](#)
- [Work with Quick Sight Topics using the Amazon Quick Sight APIs](#)

How Topics work

A Topic in Quick Sight consists of four layers:

1. **Data sources.** Your normalized datasets connect to Quick Sight from supported sources. Topics support SPICE datasets and Direct Query datasets against Amazon Redshift, Athena, Amazon S3 Tables, Snowflake, and Databricks. You cannot combine SPICE and Direct Query datasets within a single Topic.
2. **Dataset enrichment.** Each dataset is independently enriched with semantic metadata that improves natural language query accuracy. Enrichment includes column descriptions, synonyms, semantic types, calculated fields, and field exclusions. You can enrich datasets directly in Quick Sight or bring metadata in using catalogs from AWS Glue (AWS Glue Data Catalog) and Databricks Unity Catalog.
3. **Multi-dataset Topic.** The Topic serves as the unified container that brings datasets together. It holds:
 - **Datasets** (up to 12). Each dataset maintains its own enrichment and connects to its own data source.
 - **Relationships.** Join keys defined between dataset pairs tell Quick Sight how tables relate. You upload a JSON file that maps columns between datasets.

- **Custom instructions.** Persistent natural language rules that guide the AI in interpreting domain-specific terminology. These handle disambiguation, custom date logic, and business definitions.
 - **Permissions.** Owners can modify the Topic. Viewers can ask questions and use the Topic in analysis but cannot change its configuration.
4. **Consumption.** Business users interact with the Topic through multiple surfaces:
- **Chat.** Users ask natural language questions in the Amazon Quick chat interface. The LLM-powered chat agent parses the question, identifies relevant columns across datasets, constructs SQL with appropriate joins based on your defined relationships, and returns a unified answer.
 - **Analysis sheets.** Authors build visuals using fields from multiple datasets within the Topic. Quick Sight performs runtime inner joins automatically when a visual references fields from more than one dataset.

Row-level security (RLS) is enforced at the dataset level during runtime joins. Users see only the data they are authorized to access, even when queries span multiple datasets.

Creating Quick Sight Topics

Applies to: Enterprise Edition

Intended audience: Amazon Quick administrators and authors

To create a Topic, you add one or more enriched datasets, define relationships between them, and optionally add custom instructions that guide how the AI interprets cross-dataset queries.

Creating a Topic

Use the following procedure to create a Topic.

To create a Topic

1. On the Quick homepage, choose **Data** in the navigation pane at left.
2. Choose the **Topics** tab, and then choose **Create topic**.

3. Enter a topic name (for example, "Retail Sales Analytics") and an optional description, then choose **Create**.
4. Choose **Add dataset** and select the datasets you want to include in the Topic (up to 12 datasets).
5. After adding datasets, define the relationships between them. For more information, see [Defining relationships between datasets in a Topic](#).
6. (Optional) Add custom instructions that provide cross-dataset business rules. For more information, see [Adding custom instructions to a Topic](#).
7. Choose **Publish** to make the Topic available to users.

After publishing, share the Topic with business users so they can ask natural language questions or use it in analysis. For more information, see [Sharing Quick Sight Topics](#).

Prerequisites

Before creating a Topic, confirm that you have the following:

- Amazon Quick Enterprise Edition enabled with Author or Admin role.
- One or more datasets enriched with semantic metadata (column descriptions, synonyms, and semantic types). For more information, see [Data Preparation Experience \(New\)](#).
- Datasets representing a dimensional model (fact tables and dimension tables) loaded into SPICE or accessible through a supported Direct Query source.
- Permissions to create Topics and manage datasets.

Defining relationships between datasets in a Topic

Applies to: Enterprise Edition

Intended audience: Amazon Quick administrators and authors

Relationships tell Quick Sight how to join datasets when a question or visual spans multiple tables. You define relationships by specifying join keys between dataset pairs using a JSON configuration file.

The relationship graph must be a directed acyclic graph (DAG) — circular relationships are not supported. We recommend modeling your datasets in a star schema with one or more central fact tables joined to shared dimension tables.

To define relationships between datasets

1. Open the Topic that you want to configure.
2. Navigate to the **Relationships** tab.
3. Choose **Upload file** or **Create manually**.
4. Upload a JSON file defining your relationships (see the following example). After uploading, the relationship graph displays visually at left and the join key details appear at right. Select a join to verify the relationship mapping.
5. To edit a relationship, choose **Edit**, update the join columns using the dropdown selectors, and choose **Save**.

Relationship JSON format

The following example shows a star schema configuration with a central SALES_FACT table joined to dimension tables:

```
{
  "datasetPairs": [
    {
      "datasetLeft": { "datasetName": "SALES_FACT", "joinColumnNames":
["CUSTOMER_ID"] },
      "datasetRight": { "datasetName": "CUSTOMER_DIM", "joinColumnNames":
["CUSTOMER_ID"] }
    },
    {
      "datasetLeft": { "datasetName": "SALES_FACT", "joinColumnNames":
["PRODUCT_ID"] },
      "datasetRight": { "datasetName": "PRODUCT_DIM", "joinColumnNames":
["PRODUCT_ID"] }
    },
  ]
}
```

```
    "datasetLeft": { "datasetName": "SALES_FACT", "joinColumnNames": ["STORE_ID"] },
    "datasetRight": { "datasetName": "STORE_DIM", "joinColumnNames": ["STORE_ID"] }
  }
]
```

Each entry in `datasetPairs` specifies a pair of datasets and the columns used to join them. Composite keys are supported by including multiple column names in the `joinColumnNames` array.

Best practices for relationships

- **Start with a star schema.** A central fact table surrounded by dimension tables minimizes join complexity and maximizes query performance.
- **Use clean join keys.** Use integer surrogate keys where possible. Confirm matching data types on both sides. Remove null values from join key columns (nulls never match in inner joins).
- **Validate referential integrity.** Every foreign key in the fact table should exist in the dimension table.
- **Avoid circular joins.** The relationship graph must be acyclic. If your model creates a cycle, break it by removing one leg and denormalizing the redundant path.
- **Pre-join snowflake chains when feasible.** If a dimension has sub-dimensions (for example, Customer → Geography → Region), consider flattening them into a single dimension dataset to reduce join hops.

Current limitations

The following limitations apply to defined relationships:

- Relationships use *inner join* semantics for analysis sheets. Only rows with matching keys in both datasets appear in results.
- The relationship graph must be acyclic (no circular joins).
- Self-relationships (a dataset related to itself) are not supported.
- All datasets in a Topic must use the same query mode (SPICE or Direct Query).
- A Topic cannot exceed 12 datasets.

Note

When using Topics in Amazon Quick chat, the LLM-powered chat agent is not limited to inner joins. It can generate SQL with left joins, outer joins, unions, and subqueries based on your custom instructions. These join-type limitations apply only to the analysis sheet consumption path.

Adding custom instructions to a Topic

Applies to: Enterprise Edition

Intended audience: Amazon Quick administrators and authors

Custom instructions are persistent natural language rules that guide the AI engine in interpreting domain-specific terminology and cross-dataset logic. Without them, the engine interprets terms literally. For example, "this year" defaults to the calendar year. However, if your organization runs on a fiscal year starting April 1, you need an instruction to override that default.

Custom instructions are especially useful for:

- **Disambiguation rules.** When the same business term could map to multiple datasets or fields, tell the AI which one to prefer. For example: "When the user asks about 'sales', use SALES_FACT. When the user asks about 'returns' or 'refunds', use RETURNS_FACT."
- **Cross-dataset definitions.** Define metrics that span multiple datasets. For example: "Net Revenue = SUM(SALES_FACT.total_amount) - SUM(RETURNS_FACT.refund_amount)."
- **Default join behavior.** Specify the join direction that preserves intended semantics. For example: "Prefer LEFT JOIN from fact tables to dimension tables so that facts without matching dimension records are not silently dropped."
- **Custom date logic.** For example: "Fiscal year starts April 1. Interpret 'this year' using fiscal year boundaries."

To add custom instructions

1. Open the Topic that you want to configure.

2. Navigate to the **Custom instructions** tab.
3. Choose **Edit**.
4. Enter your instructions as natural language rules. Use bullet-point style for clarity.
5. Choose **Save changes**.

Example custom instructions

The following example shows a production-ready set of topic-level instructions for a retail analytics Topic:

Disambiguation:

- "sales", "revenue", "orders" -> use SALES_FACT
- "returns", "refunds" -> use RETURNS_FACT
- "net sales", "net revenue" -> join SALES_FACT LEFT JOIN RETURNS_FACT on order_line_id

Cross-dataset metrics:

- Net Revenue = SUM(SALES_FACT.total_amount) - SUM(RETURNS_FACT.refund_amount)
- Return Rate = COUNT(RETURNS_FACT.return_id) / COUNT(SALES_FACT.order_line_id)

Default joins:

- SALES_FACT LEFT JOIN CUSTOMER_DIM on customer_id
- SALES_FACT LEFT JOIN PRODUCT_DIM on product_id

Date handling:

- Fiscal year starts April 1. Interpret "this year" using fiscal year boundaries.
- YTD: current fiscal year up to and including today.

Best practices for custom instructions

- Keep instructions concise. Prefer bullet-point rule lists over prose paragraphs.
- Use topic-level instructions for cross-dataset logic only. Single-dataset semantics (grain, keys, aggregation rules) belong in the dataset's own enrichment metadata.
- Avoid contradictions between dataset-level and topic-level instructions.
- Test your instructions by asking questions in chat and reviewing the generated SQL in the Explanation panel.

Using Topics in Quick Sight analysis

Applies to: Enterprise Edition

Intended audience: Amazon Quick administrators and authors

You can use a Topic as the data model for your Quick Sight analysis. When you create an analysis from a Topic, you can select fields from any of the datasets in the Topic. Quick Sight automatically performs runtime inner joins based on the relationships you defined, assembling precisely the join it needs for each visual.

This means you no longer need to pre-join all your tables into a single flat dataset before building an analysis. Each dataset maintains its own level of detail, and Quick Sight joins only the relevant tables at analysis time.

Creating an analysis from a Topic

To create an analysis from a Topic

1. Navigate to the Topic and choose **Create analysis**.
2. Select **Interactive sheet** and choose **Create**.
3. In the analysis, select fields from multiple datasets and add them to a visual. Quick Sight performs runtime joins automatically based on the defined relationships.
4. (Optional) Create calculated fields that reference columns from different datasets.

How runtime joins work in analysis

When a visual references fields from multiple datasets:

1. Quick Sight identifies which datasets contain the fields used in the visual.
2. It uses the defined relationships to determine the join path between datasets.
3. It performs inner joins at runtime using the specified join keys.
4. The result is displayed in the visual with the correct aggregation and grouping.

Note

Runtime joins in analysis use inner join semantics. Only rows with matching keys in both datasets appear in results. If you need outer join behavior for analytical scenarios, use the Topic in Amazon Quick chat instead, where the chat agent can generate SQL with outer joins based on your custom instructions.

Benefits of using Topics in analysis

- **Less upfront data preparation.** Define relationships once. Quick Sight joins only the relevant tables at analysis time.
- **Preserved native granularity.** Each dataset maintains its own level of detail, avoiding measure duplication across grains.
- **Reuse across analyses.** A single Topic with defined relationships serves multiple analytical use cases without rebuilding datasets.
- **Independent refresh schedules.** Each dataset can be refreshed at different cadences (hourly, daily, monthly) based on data volatility.
- **Row-level security at runtime.** RLS rules are enforced during runtime joins, so data-access policies apply consistently across datasets.

Using Topics in Amazon Quick chat

Applies to: Enterprise Edition

You can use Topics in Amazon Quick chat to ask natural language questions that span multiple datasets. The LLM-powered chat agent uses your defined relationships, dataset enrichment metadata, and custom instructions to generate cross-dataset SQL queries and return unified answers.

How cross-dataset chat queries work

When you ask a natural language question against a Topic, the chat agent performs the following steps:

1. **Intent parsing.** The agent identifies which columns map to your terms by matching against column names, descriptions, and synonyms from the enrichment metadata. It determines which datasets contain the relevant measures and dimensions.
2. **Relationship traversal.** Using the defined join keys and custom instructions, the agent determines the join path between identified datasets. It can traverse relationships to connect fact tables to the necessary dimension tables.
3. **SQL generation.** The agent constructs a SQL query with appropriate JOIN clauses, aggregation, and GROUP BY for the requested dimensions.
4. **Result presentation.** The answer is returned as a visualization or table, with the generated SQL available for inspection in the Explanation panel.

Chat capabilities with Topics

The LLM-powered chat agent supports richer SQL generation than the defined relationship path used in analysis sheets. When you configure a Topic for chat with custom instructions, the agent can generate:

- Inner, left, right, and full outer joins
- Union queries across tables with the same schema
- Subqueries for negation patterns (for example, "customers who have never ordered")
- Cross-grain comparisons (for example, daily actuals versus monthly targets)
- Self-joins for recursive hierarchies (for example, employee-manager relationships)

The richer your enrichment metadata and custom instructions, the more accurately the agent interprets ambiguous questions.

Starting a chat with a Topic

To chat with a Topic:

- Navigate to the Topic and choose the chat icon. Your Topic is automatically available as the context.
- Or, from Amazon Quick chat, use the data filter to select **Specific data and apps** and choose your Topic.

Ask natural language questions that span multiple datasets. For example:

- "Show total sales by customer segment and store region"
- "What is the return rate by product category?"
- "Which stores are below 80% of their monthly target?"

Use the Explanation panel to view the generated SQL and verify that the correct datasets and joins were used.

Differences from legacy Topics in chat

The chat experience differs significantly between new Topics and legacy Topics:

Aspect	New Topics	Legacy Topics
AI model	LLM-powered chat agent	Legacy ML-based fuzzy search model
Cross-dataset queries	Yes — traverses relationships and generates cross-dataset SQL with runtime joins	No — selects one dataset and queries only that dataset
Join types	Inner, left, right, full outer, union, subquery	Not applicable (single dataset)
SQL visibility	Generated SQL shown in Explanation panel	Not available

 Note

You can chat with both new Topics and legacy Topics. However, legacy Topics use the legacy ML-based model and do not support cross-dataset queries. For information about legacy Topics, see [Working with legacy Topics](#).

Sharing Quick Sight Topics

Applies to: Enterprise Edition

Intended audience: Amazon Quick administrators and authors

After you create and publish a Topic, share it with others in your organization. Sharing a Topic allows your users to ask questions in Amazon Quick chat and use the Topic as a data model in analysis sheets.

To share a Topic

1. From the Topic page, select the ellipsis menu and choose **Share**.
2. Search for and add specific users or groups.
3. Set permission levels (**Owner** or **Viewer**) and choose **Done**.

Permission Level	Can Ask Questions	Can Modify Topic	Can Use in Analysis
Owner	Yes	Yes	Yes
Viewer	Yes	No	Yes

Quick Sight enforces row-level security (RLS) and column-level security (CLS) at the dataset level. Access controls are preserved through the Topic's semantic layer, regardless of how users access the Topic.

Working with legacy Topics

Applies to: Enterprise Edition

Intended audience: Amazon Quick administrators and authors

If you created Topics before the multi-dataset Topics feature launched, your existing Topics are now classified as *legacy Topics*. Legacy Topics continue to work as before — their behavior and capabilities are unchanged.

Differences between legacy Topics and new Topics

Feature	New Topics	Legacy Topics
Purpose	Multi-dataset semantic layer with cross-dataset relationships, runtime joins for both chat and analysis	Single-dataset semantic layer for natural language Q&A
Datasets	Up to 12 datasets with defined relationships	Multiple datasets can be added, but no cross-dataset joins at query time
Chat behavior	LLM-powered chat agent generates cross-dataset SQL with runtime joins	ML-based fuzzy search model selects one dataset, then queries only that dataset
Analysis support	Use as data model for analysis sheets with runtime joins	Limited to NLQ search bar in analysis
Metadata location	Dataset enrichment lives in the dataset itself; Topic holds cross-dataset logic	Metadata (synonyms, calculated fields, named entities, filters) stored in the separate legacy Topic object
Relationships	Explicit join keys defined between dataset pairs	Not supported

Migrating from legacy Topics

You can migrate your business context from legacy Topics into enriched datasets using Dataset Enrichment in the new data preparation experience. This moves column descriptions, synonyms, calculated fields, named entities, filters, and custom instructions from the separate legacy Topic object into the dataset itself.

After migrating your dataset-intrinsic metadata into Dataset Enrichment, you can create a new multi-dataset Topic that uses the enriched datasets with defined relationships.

For detailed migration guidance, see [Data Preparation Experience \(New\)](#).

Navigating legacy Topics

From your Quick start page, you can view and manage legacy Topics by selecting **Data** in the navigation pane at left, then choosing the **Topics** tab. Legacy Topics are labeled as such in the topics list.

You can also access a legacy Topic from within an Amazon Quick Sight analysis by choosing the topic icon in the top navigation bar.

Creating a legacy Topic

Note

New Topics created going forward use the multi-dataset Topic experience. The following information applies to legacy Topics that were created before this feature launched.

Legacy Topics are collections of one or more datasets that represent a subject area that your business users can ask questions about. When you create a legacy Topic, Quick Sight uses ML-powered automated data prep to help select fields, assign friendly names and synonyms, and detect semantic types.

There are two ways to create a legacy Topic:

- Create the Topic by selecting a dataset from the Topics page.
- Create the Topic using an analysis, where automated data prep learns from how you analyze your data.

Legacy Topic workspace

After you open a legacy Topic, the workspace displays four tabs:

- **Summary** — Shows suggestions for improving the Topic, engagement metrics and KPIs, and the datasets included in the Topic.
- **Data** — Shows all the fields included in the Topic. Configure field metadata (names, descriptions, synonyms, semantic types, aggregations) to make the Topic natural-language-friendly.
- **User Activity** — Shows all questions received and the overall feedback for each question.
- **Verified Answers** — Shows questions that you have preconfigured visuals for.

Working with datasets in a legacy Topic

You can add additional datasets to a legacy Topic, remove datasets, set refresh schedules, and add calculated fields, filters, and named entities.

- **Adding datasets** — On the Summary page, choose Data, then under Datasets, choose Add datasets.
- **Removing datasets** — On the Summary page, choose Data, then under Datasets, choose the dataset menu and select Remove from topic.
- **Refreshing datasets** — You can set a refresh schedule or manually refresh datasets in a legacy Topic. For SPICE datasets, you can sync with the SPICE refresh schedule.
- **Calculated fields** — Create new fields using combinations of existing fields with supported functions.
- **Filters** — Add filters to support questions that map to multiple values in the data.
- **Named entities** — Create collections of fields that display together in an answer (for example, an "Address" entity combining Branch Name, State, and City).

Making legacy Topics natural-language-friendly

To improve the accuracy of answers in a legacy Topic, configure the following metadata:

1. **Dataset names and descriptions** — Give datasets friendly names so readers can identify them.
2. **Date fields** — Specify default dates and time basis for each dataset.
3. **Exclude unused fields** — Remove irrelevant fields from the Topic to improve accuracy.
4. **Rename fields** — Give fields user-friendly names and descriptions.
5. **Synonyms** — Add alternative names that readers might use to refer to fields or values.
6. **Semantic types** — Specify the type of information in each field (location, currency, date, person, etc.).
7. **Field roles and aggregations** — Specify whether fields are dimensions or measures, and set default aggregations.

Sharing legacy Topics

You can share legacy Topics with other users by assigning them as Viewers or Co-owners. Viewers can see the Topic in the search bar and ask questions. Co-owners can also change the Topic configuration.

Reviewing legacy Topic performance

After sharing a legacy Topic, you can review its performance on the Summary and User Activity tabs. This shows how many questions were asked, feedback received, and any questions that were unanswerable or disambiguated.

Refreshing legacy Topic indexes

When you create a legacy Topic, Quick Sight creates and maintains an index of unique string values for fields included in the Topic. This index is used to generate correct answers, provide autocomplete suggestions, and suggest mappings of terms to columns or data values.

To refresh a legacy Topic index, refresh the datasets in the Topic. You can manually refresh all datasets or set a recurring refresh schedule.

Work with Quick Sight Topics using the Amazon Quick Sight APIs

Applies to: Enterprise Edition

Intended audience: Amazon Quick developers

Use this section to learn how to work with Quick Sight Topics using the Amazon Quick Sight command line interface (CLI).

Prerequisites

Before you begin, make sure that you have an AWS Identity and Access Management (IAM) role that grants the CLI user access to call the Quick Sight API operations. The following table shows which permissions must be added to the IAM policy to use specific API operations.

API operation	IAM policy
CreateTopic	quicksight:CreateTopic quicksight:PassDataSet
ListTopics	quicksight:ListTopics

API operation	IAM policy
DescribeTopic	quicksight:DescribeTopic
UpdateTopic	quicksight:UpdateTopic quicksight:PassDataSet
DeleteTopic	quicksight>DeleteTopic
UpdateTopicPermissions	quicksight:UpdateTopicPermissions
DescribeTopicPermissions	quicksight:DescribeTopicPermissions
CreateTopicRefreshSchedule	quicksight:CreateTopicRefreshSchedule
ListTopicRefreshSchedules	quicksight:ListTopicRefreshSchedules
DescribeTopicRefreshSchedule	quicksight:DescribeTopicRefreshSchedule
UpdateTopicRefreshSchedule	quicksight:UpdateTopicRefreshSchedule
DeleteTopicRefreshSchedule	quicksight>DeleteTopicRefreshSchedule

The following example creates a new Topic.

```
aws quicksight create-topic
--aws-account-id AWSACCOUNTID
--topic-id TOPICID
--topic TOPIC
```

The following example updates a Topic.

```
aws quicksight update-topic  
--aws-account-id AWSACCOUNTID  
--topic-id TOPICID  
--topic TOPIC
```

The following example provides a list of all Topics in an account.

```
aws quicksight list-topics  
--aws-account-id AWSACCOUNTID
```

The following example deletes a Topic.

```
aws quicksight delete-topic  
--aws-account-id AWSACCOUNTID  
--topic-id TOPICID
```

The following example provides information about how a Topic was configured.

```
aws quicksight describe-topic  
--aws-account-id AWSACCOUNTID  
--topic-id TOPICID
```

Working with data stories in Amazon Quick Sight

With Generative BI with Quick Sight, authors and readers can generate a first draft of their data story quickly. Use prompts and visuals to produce a draft that incorporates the details that you provide. Data story drafts are not meant to replace your own ideas or to perform analysis. Rather, data stories are a starting point to customize and expand on as needed. The contextual recommendations and suggestions combine your prompt with selected visuals to provide relevant details that are tailored to your data story. For more information about this, see [Generative BI with Quick Sight](#).

Use the following topics to create, modify, and share a data story.

Topics

- [Creating a data story with Generative BI](#)
- [Personalize data stories in Amazon Quick Sight](#)
- [Viewing a generated data story in Amazon Quick Sight](#)

- [Editing a generated data story in Amazon Quick Sight](#)
- [Adding themes and animations to a data story in Amazon Quick Sight](#)
- [Sharing a data story in Amazon Quick Sight](#)

Creating a data story with Generative BI

Use the following procedure to create a data story with Generative BI.

To create a data story

1. Open the [Quick console](#).
2. At left, choose **Stories**.
3. On the **Data stories** page, choose **New Data Story**.
4. In the **Story** screen that appears, navigate to the **Build story** modal and input a data story prompt that you would like to generate. For the best results, don't phrase the prompt like a question. Instead, type the data story that you want Quick Sight to build. For example, say you want to create a data story about the most commonly performed medical procedures by region. A good prompt for this use case is "Build a data story about most commonly performed procedures by physicians in various regions. Also, show the specialties where patients are admitted the most. Recommend where we need to staff more physicians by specialty, and include at least four points of supporting data."

You can optionally skip this step and manually create your data story. If you choose to forego entering a prompt, you still need to add a visual to the data story.

5. Under **Select visuals**, choose **Add**.
6. Choose the dashboard that contains the visuals that you want to use, and then choose the visuals that you want. You can add up to 20 visuals to a data story.

If you don't see the dashboard that you want to use, use the **Find your dashboards** search bar at the top of the modal.

You can choose visuals from any number of dashboards that you have sharing permissions to. Visuals that show a **Restricted** badge have permissions that restrict them from being added to a data story. A visual might be restricted for one of the following reasons:

- The dataset is connected to a data source that uses trusted identity propagation with Amazon Redshift.

- The dataset is located inside of a restricted folder.
7. (Optional) Use the **Select documents** section to upload up to 5 documents to be used in the data story. Each document can't exceed 10MB. These documents are only used to generate the data story and are not stored in Quick Sight. The following image shows the **Select documents** section of the **Build story** screen.
 8. (Optional) If your Quick account is connected to an Amazon Q Business application, check the **Use insights from Amazon Q Business** checkbox to augment your data story with unstructured data sources from Amazon Q Business. For more information about connecting a Quick account to a Amazon Q Business application, see [Augmenting Amazon Quick Sight insights with Amazon Q Business](#).
 9. Choose **Build**.

After the data story generates, review the data story and choose from the following options:

- **Keep** – Saves the generated content to the canvas. When you choose this option, the **Build story** modal closes and you can start editing your data story.
- **Try again** – Allows users to edit the prompt and generate a new data story.
- **Discard** – Deletes the generated data story.

Personalize data stories in Amazon Quick Sight

User location and job-related information from your IAM Identity Center instance are leveraged to generate personalized data stories that are more relevant to authors and readers. For example, when an author in the US issues the prompt "Write a business strategy focusing on a plan on how to increase the revenue in my location", insights related to the US in the data story's narrative are automatically included. If the author wants the data story to focus on another country such as Canada, they can specify this in the prompt.

For personalization to work, you must add country and job title for users in the IAM Identity Center instance that is connected to your Quick account. For more information, see [Add users to your IAM Identity Center directory](#) in the IAM Identity Center User Guide.

User data in your IAM Identity Center instance is connected to your application environment by default. This means that all data stories are personalized by default. You can choose to [opt out of personalization](#) at any time in the Account settings menu in the QuickSight administration console.

Note

Personalization in data stories is currently available in the US East (N. Virginia) and US West (Oregon) AWS Regions.

Viewing a generated data story in Amazon Quick Sight

After you generate and keep a data story, you can access that data story from the **Data stories** page. To view a data story, choose the data story that you want to view to open the story editor.

As you create and modify a data story, you can preview how the data story looks to readers. To preview a generated data story, choose the **Preview** icon at the top of the page. To exit the preview, choose **BACK TO EDITOR**.

Editing a generated data story in Amazon Quick Sight

After you create and keep a data story, you can modify its content to better fit your needs. You can format data story text, add images, edit visuals, and add new blocks.

Stories are made up of different *blocks* that act as containers for text, visuals, and images that you want to include in your data story. Each block can be formatted independently from other blocks in the data story, similar to the sections of a pixel perfect report.

To format the text of a data story, use the toolbar at the top of the page. The toolbar offers font settings so you can customize the font type, style, color, size, spacing, text highlights, and alignment. You can also use the toolbar to add columns to a data story block.

Use one of the following options to add a visual to a data story.

- Use the **Visuals** pane to drag and drop a visual into a data story. Only the visuals that you chose when you generated the data story are shown in the **Visuals** pane.

You can also choose the **Add (+)** icon in the **Visuals** pane to add new visuals that can be dragged and dropped into the data story. Each data story can contain up to 20 visuals.

- Choose the data story block that you want to add an image to. When a cursor appears, enter a forward slash ("/") to insert an image or visual to that data story block.

To edit a visual in a data story, choose the visual that you want to change, and then choose the **Properties** icon. In the properties pane that appears, you can perform the following actions:

- Change, hide, or show the visual's title. By default, the visual title is displayed.
- Change, hide, or show the visual's subtitle. By default, the visual subtitle is hidden.
- Hide or show data labels. By default, data labels are hidden.
- Hide, show, or change the position of the legend. By default, the legend is hidden.

To add a new block to a data story, choose the plus (+) icon at the bottom of any existing block. Then choose the layout option that you want. You can also move, duplicate, or delete a block from the **Block options** (three dots) icon at the top of each block.

To change the layout of items in a block, you can drag and drop the items wherever you want with the six-dot icon next to each item.

Adding themes and animations to a data story in Amazon Quick Sight

You can add themes and animations to the stories that you generate. To add a theme or animation to a data story, choose the **Story style** icon.

In the **Story style** pane that appears, you can perform the following actions:

- For **THEMES**, choose a theme that you think best fits your data story.
- For **ANIMATIONS**, choose an animation style and speed. For animation types, you can choose **None**, **Fade**, or **Slide**. The default animation is **None**. For **Speed**, choose **Slow**, **Medium**, or **Fast**. The default speed is **Medium**.

Sharing a data story in Amazon Quick Sight

Use the following procedure to share a data story.

To share a data story

1. In the story editor of the data story that you want to share, choose the **Share** icon at the top right.

Alternatively, you can choose the **Share** icon at the top of a data story preview.

2. In the **Share data story** modal that appears, enter the users or groups that you want to share the data story with.
3. (Optional) To save a link for the published data story to your clipboard, choose **Copy Link**.
4. Choose **Publish & Share**.

If you try to share a story and receive a message that the story cannot be shared, contact the owner of the dashboard ask them to toggle on the **Allow sharing data stories** switch. For more information about this switch, see [Tutorial: Create an Amazon Quick Sight dashboard](#).

If you try to share a data story and receive an error message, contact the owner of the dashboard or your Quick account admin for assistance.

After you share a data story, users you shared the story with receive a notification email with a link to the story. You can access the data story from the **Data stories** page of their Quick accounts. You can also share the copied link to the data story with users that can access the data story.

You can't share a data story that contains restricted data. If you try to share a story that contains restricted data, an error message appears that lists all restricted visuals that are a part of the story. If desired, remove the restricted visuals from your data story before sharing it with users.

When you edit a published data story, republish the data story for the changes to propagate to your end users.

Working with scenarios in Amazon Quick Sight

Quick users with Admin Pro, Author Pro, or Reader Pro roles can use scenarios to analyze complex business problems with simple natural language.

To get started with scenarios, a Quick user describes a problem that they want to solve and adds relevant data from Quick Sight or from their computer to be used in the data analysis. Alternatively, users can let Amazon Q search for all relevant data that can be used to solve the problem. Amazon Q returns a series of analyses or prompts to dive deeper into the data. Users can also enter their own prompts to create a custom analysis. After a new prompt is received, Amazon Q breaks down the analysis into steps and executes them. Outputs include specific data insights, interactive visuals, and an analysis of what the findings might mean for the business with suggested next actions.

Scenarios can help Quick Pro users to perform the following tasks:

- Automate tedious, error-prone, and inefficient manual data tasks
- Modify, extend, or reuse past analyses to quickly adapt to business changes
- Dive deeper into data than spreadsheets allow

Use the following topics to create and work with scenarios in Amazon Quick Sight.

Topics

- [Considerations for Quick Sight scenarios](#)
- [Creating an Amazon Quick Sight scenario](#)
- [Working with threads in an Amazon Quick Sight scenario](#)
- [Working with data in an Amazon Quick Sight scenario](#)

Considerations for Quick Sight scenarios

The following considerations apply to Amazon Quick Sight scenarios.

- Amazon Quick Sight scenarios are available to users that have Admin Pro, Author Pro, or Reader Pro roles in Amazon Quick. For information about updating a user to a Quick Pro role, see [Get started with Generative BI](#).
- Scenarios are available in specific AWS Regions listed in [Supported AWS Regions for Amazon Q in Quick](#).

After you review the considerations for Quick Sight scenarios, see [Creating an Amazon Quick Sight scenario](#) to get started with scenarios in Amazon Quick Sight.

Creating an Amazon Quick Sight scenario

Amazon Quick Pro users can create scenarios from Quick Sight dashboards, or from the **Scenarios** section on the Quick Sight home page. Users can create as many scenarios as they need. Each user can have up to 3 active scenarios at a time. Each Quick Sight account supports up to 10 active scenarios at a time. Use the following procedure to create a scenario in Amazon Quick Sight.

Create a new scenario

1. Open the [Quick console](#).
2. Perform one of the following actions:
 - a. Open any dashboard, and look for one of the following:
 - Choose **Analyze this dashboard in a Scenario**, if available, at the top of the dashboard.
 - From a visual on the dashboard, open the drop-down menu and choose **Explore scenario**.
 - Choose **Build**, and then choose **Scenario**.

- b. On the Quick home page, choose **Scenarios**. On the **Scenarios**, choose **New Scenario**.
3. The new scenario appears. In the text box, describe the problem that you want to solve. This input is the starting point for all of the data pivots and manipulations that will occur in the scenario. The description that you provide can be as broad or as specific as you want, for example "analyze usage trends" or "compute month-over-month and year-over-year changes in usage based on last month's data."
4. Add the data that you want to use in the scenario. You can choose data from Quick Sight dashboards, or you can upload files from your computer. When you choose data from a dashboard, a preview of the selected data is generated for you to review. For more information about previewing and editing data in Quick Sight scenarios, see [Working with data in an Amazon Quick Sight scenario](#).

The following limits apply to the data that is used in a scenario:

- You can add up to 10 data sources to a scenario.
- Up to 20 visuals can be selected from a dashboard at a time.
- Uploaded files must be in .xlsx or .csv format and can't exceed 1 GB.
- Data sources can have up to 200 columns.

If you don't add data to the scenario, Amazon Q automatically searches your Quick Sight dashboards to find data related to your problem statement from the previous step.

5. Choose **Start analysis**.

When you start an analysis in a Quick Sight scenario, Quick Sight prepares your data for analysis and returns a new *thread*. The thread contains generated prompts that can be used to solve the problem that you described in the scenario. A thread is a turn based contextual conversation that consists of user prompts and Amazon Q responses that you can use to drill down on a specific scenario. You can use threads to write prompts that assume that Amazon Q remembers what was previously discussed in the thread. You can choose a prompt to continue the thread, or you can choose the plus sign (+) above the thread to start a new thread. The new thread uses a different prompt than the first thread that you created. For more information about working with threads, see [Working with threads in an Amazon Quick Sight scenario](#).

Working with threads in an Amazon Quick Sight scenario

After you create a scenario in Quick Sight, the data that Amazon Q generates is presented in *threads* and *blocks*. A thread is a vertical chain of prompts and responses. A block is a single prompt and response pair. Each thread can contain up to 15 blocks, and each scenario can contain up to 50 blocks total across multiple threads.

When a new thread is created, a list of Amazon Q-generated prompts appears inside of a new block. When you choose one of the prompts to drill down on, Amazon Q analyzes the data that is relevant to the chosen prompt and returns a summary of all data findings, forecasts, and conclusions that can be drawn from the analysis.

To continue the thread and dive deeper into the prompt, choose the plus sign (+) located below the block to create a new block that contains a new list of generated prompts that factor in the findings from the previous block. To start a new thread that analyzes a different aspect of the data, choose the plus sign (+) above any block in the scenario to create a new thread.

Blocks can be collapsed, duplicated, or deleted from a scenario, as long as the block that you want to change has finished loading. Use the following procedures to make changes to a scenario block.

To collapse, duplicate, or delete a block

1. Open the [Quick console](#).
2. Choose **Scenarios** from the options pane, and then choose the scenario that you want to change.
3. Navigate to the block that you want to change and choose the ellipsis (...) in the top right of the block.
4. Perform one of the following actions:
 - To collapse the block, choose **Collapse**. To expand a collapsed block, choose the ellipsis in the top right of the block, and then choose **Expand**.
 - To duplicate the block, choose **Duplicate**. The block is duplicated and placed in a new thread next to the original block.
 - To delete the block, choose **Delete**.

You can also modify the prompt of a block to better match your use case. Use the following procedure to modify a block prompt.

To modify the prompt of a block

1. Open the [Quick console](#).
2. Choose **Scenarios** from the options pane, and then choose the scenario that you want to change.
3. Navigate to the block that you want to change and choose **Modify block**.
4. In the **Modify block** popup that appears, enter a new description for the block, and then choose **Apply**.

After you modify a prompt, Amazon Q analyzes the data and returns a new generated analysis that reflects the changes that were made to the prompt.

Working with data in an Amazon Quick Sight scenario

When you create a scenario in Amazon Quick Sight, you can preview and modify the data that the scenario uses to generate summaries. Use the following sections to learn about the ways Quick users can interact with data in a scenario.

Topics

- [Adding more data to a scenario](#)
- [Editing data in a preview](#)
- [Editing data in a snapshot](#)

Adding more data to a scenario

After you create a scenario in Amazon Quick Sight, you can add more data to the scenario at any time. Use the following procedure to add data to an Amazon Quick Sight scenario.

To add data to an existing Amazon Quick Sight scenario

1. Open the [Quick console](#).
2. Choose **Scenarios** from the options pane, and then choose the scenario that you want to add more data to.
3. Choose the **Data Source** icon in the actions bar to open the **Data** pane.
4. Perform one of the following actions:

- a. To add Quick Sight data to the scenario, choose **Find Data**, and then choose the dataset or dashboard visuals that you want to add to the scenario. After you have selected all of the Quick Sight data that you want to add to the scenario, choose **Add**.
- b. To upload a file from your computer to the scenario, choose **Upload File**.

The following limits apply to the data that is added to a scenario:

- You can add up to 10 data sources to a scenario.
- Up to 20 visuals can be selected from a dashboard at a time.
- Uploaded files must be in .xlsx or .csv format and can't exceed 1 GB.
- Data sources can have up to 200 columns.

After you add new data to a scenario, Amazon Q includes the data in all new analyses.

Editing data in a preview

When you choose data from a Quick Sight dashboard to be used in a scenario, a preview of the data is generated for review before it's added to the analysis. If needed, the following changes can be made to dashboard data in the preview state:

- **Filters** – If you only want to analyze a subset of the available data or if you need to reduce the number of rows that are included in the scenario, you can apply filters to the data.
- **Sort** – If the available data exceeds 1 million rows and you want to prioritize the retention of the values in a specific column, you can sort the data to fit your needs.

Editing data in a snapshot

When you add dashboard or external data to a scenario, Quick Sight creates a snapshot of the data sources to be reviewed. To see a snapshot of the data used in a scenario, choose the **Data Source** icon in the actions bar. This opens the **Data** pane, and then you can choose the data snapshot that you want to review.

You can perform the following actions on a data snapshot:

- To update the title of the data snapshot, choose the pencil icon next to the title and enter a new title for the snapshot.

- Choose the **Filter** icon to filter the data that is used in the scenario. This option can be used if you want the scenario to only use a subset of the data that is added to the scenario.
- Choose the **Sort** icon to sort the data that is used in the scenario. This option can be used to prioritize the retention of specific columns if the data exceeds 1 million rows.
- Choose the **Fields list** icon to choose which fields are included in the scenario. This option can be used to control which columns are used in the scenario.

When you are finished updating the scenario data, close the **Data** pane.

Troubleshooting Amazon Quick Sight

Use this information to help you diagnose and fix common issues that you can encounter when using Amazon Quick Sight.

Note

Need more help? You can visit the Amazon Quick Sight [User Community](#) or the [AWS forums](#). See also the [Quick Sight Resource Library](#).

Topics

- [Resolving Amazon Quick Sight issues and error messages](#)
- [Connectivity issues when using Amazon Athena with Amazon Quick Sight](#)
- [Data source connectivity issues for Amazon Quick Sight](#)
- [Login issues with Quick Sight](#)
- [Visual issues with Quick Sight](#)

Resolving Amazon Quick Sight issues and error messages

If you are having difficulties or receiving an error message, there's a few ways that you can go about resolving the issue. Following are some resources that can help:

- For errors during dataset ingestion (importing data), see [SPICE ingestion error codes](#).
- For technical user questions, visit the [User Community](#).
- For administrator questions, see the [AWS forums](#).

- If you need more customized assistance, contact AWS Support. To do this while you are signed in to your AWS account, choose **Support** at upper right, and then choose **Support Center**.

Connectivity issues when using Amazon Athena with Amazon Quick Sight

Following, you can find information about troubleshooting issues that you might encounter when using Amazon Athena with Amazon Quick Sight.

Before you try troubleshooting anything else for Athena, make sure that you can connect to Athena. For information about troubleshooting Athena connection issues, see [I can't connect to Amazon Athena](#).

If you can connect but have other issues, it can be useful to run your query in the Athena console (<https://console.aws.amazon.com/athena/>) before adding your query to Amazon Quick Sight. For additional troubleshooting information, see [Troubleshooting](#) in the *Athena User Guide*.

Topics

- [Column not found when using Athena with Amazon Quick Sight](#)
- [Invalid data when using Athena with Amazon Quick Sight](#)
- [Query timeout when using Athena with Amazon Quick Sight](#)
- [Staging bucket no longer exists when using Athena with Amazon Quick Sight](#)
- [Table incompatible when using AWS Glue with Athena in Amazon Quick Sight](#)
- [Table not found when using Athena with Amazon Quick Sight](#)
- [Workgroup and output errors when using Athena with Quick Sight](#)

Column not found when using Athena with Amazon Quick Sight

You can receive a "column not found" error if the columns in an analysis are missing from the Athena data source.

In Amazon Quick Sight, open your analysis. On the **Visualize** tab, choose **Choose dataset**, **Edit analysis data sets**.

On the **Data sets in this analysis** screen, choose **Edit** near your dataset to refresh the dataset. Amazon Quick Sight caches the schema for two minutes. So it can take two minutes before the latest changes display.

To investigate how the column was lost in the first place, you can go to the Athena console (<https://console.aws.amazon.com/athena/>) and check the query history to find queries that edited the table.

If this error happened when you were editing a custom SQL query in preview, verify that the name of the column in the query, and check for any other syntax errors. For example, check that the column name isn't enclosed in single quotation marks, which are reserved for strings.

If you still have the issue, verify that your tables, columns, and queries comply with Athena requirements. For more information, see [Names for Tables, Databases, and Columns](#) and [Troubleshooting](#) in the *Athena User Guide*.

Invalid data when using Athena with Amazon Quick Sight

An invalid data error can occur when you use any operator or function in a calculated field. To address this, verify that the data in the table is consistent with the format that you supplied to the function.

For example, suppose that you are using the function `parseDate(expression, ['format'], ['time_zone'])` as `parseDate(date_column, 'MM/dd/yyyy')`. In this case, all values in `date_column` must conform to 'MM/dd/yyyy' format ('05/12/2016'). Any value that isn't in this format ('2016/12/05') can cause an error.

Query timeout when using Athena with Amazon Quick Sight

If your query times out, you can try these options to resolve your problem.

If the failure was generated while working on an analysis, remember that the Amazon Quick Sight timeout for generating any visual is two minutes. If you're using a custom SQL query, you can simplify your query to optimize running time.

If you are in direct query mode (not using SPICE), you can try importing your data to SPICE. However, if your query exceeds the Athena 30-minute timeout, you might get another timeout while importing data into SPICE. For the most current information on Athena limits, see [Amazon Athena Limits](#) in the *AWS General Reference*.

Staging bucket no longer exists when using Athena with Amazon Quick Sight

Use this section to help solve this error: **"The staging bucket for this query result no longer exists in the underlying data source."**

When you create a dataset using Athena, Amazon Quick Sight creates an Amazon S3 bucket. By default, this bucket has a name similar to "aws-athena-query-results-*<REGION>*-*<ACCOUNTID>*". If you remove this bucket, then your next Athena query might fail with an error saying the staging bucket no longer exists.

To fix this error, create a new bucket with the same name in the correct AWS Region.

Table incompatible when using AWS Glue with Athena in Amazon Quick Sight

If you are getting errors when using AWS Glue tables in Athena with Amazon Quick Sight, it might be because you're missing some metadata. Follow these steps to find out if your tables don't have the `TableType` attribute that Amazon Quick Sight needs for the Athena connector to work. Usually, the metadata for these tables wasn't migrated to the AWS Glue Data Catalog. For more information, see [Upgrading to the AWS Glue Data Catalog Step-by-Step](#) in the *AWS Glue Developer Guide*.

If you don't want to migrate to the AWS Glue Data Catalog at this time, you have two options. You can recreate each AWS Glue table through the AWS Glue Management Console. Or you can use the AWS CLI scripts listed in the following procedure to identify and update tables with missing `TableType` attributes.

If you prefer to use the CLI to do this, use the following procedure to help you design your scripts.

To use the CLI to design scripts

1. Use the CLI to learn which AWS Glue tables have no `TableType` attributes.

```
aws glue get-tables --database-name <your_database_name>;
```

For example, you can run the following command in the CLI.

```
aws glue get-table --database-name "test_database" --name  
"table_missing_table_type"
```

Following is a sample of what the output looks like. You can see that the table "table_missing_table_type" doesn't have the `TableType` attribute declared.

```
{  
  "TableList": [  
    {
```



```
"Retention": 0,
"UpdateTime": 1522368588.0,
"PartitionKeys": [
  {
    "Name": "year",
    "Type": "string"
  },
  {
    "Name": "month",
    "Type": "string"
  },
  {
    "Name": "day",
    "Type": "string"
  }
],
"LastAccessTime": 1513804142.0,
"Owner": "owner",
"Name": "table_missing_table_type",
"Parameters": {
  "delimiter": ",",
  "compressionType": "none",
  "skip.header.line.count": "1",
  "sizeKey": "75",
  "averageRecordSize": "7",
  "classification": "csv",
  "objectCount": "1",
  "typeOfData": "file",
  "CrawlerSchemaDeserializerVersion": "1.0",
  "CrawlerSchemaSerializerVersion": "1.0",
  "UPDATED_BY_CRAWLER": "crawl_date_table",
  "recordCount": "9",
  "columnsOrdered": "true"
},
"StorageDescriptor": {
  "OutputFormat": "org.apache.hadoop.hive ql.io.HiveIgnoreKeyTextOutputFormat",
  "SortColumns": [],
  "StoredAsSubDirectories": false,
  "Columns": [
    {
      "Name": "col1",
      "Type": "string"
    }
  ]
}
```

```

        "Name": "col2",
        "Type": "bigint"
    }
],
"Location": "s3://myAthenatest/test_dataset/",
"NumberOfBuckets": -1,
"Parameters": {
    "delimiter": ",",
    "compressionType": "none",
    "skip.header.line.count": "1",
    "columnsOrdered": "true",
    "sizeKey": "75",
    "averageRecordSize": "7",
    "classification": "csv",
    "objectCount": "1",
    "typeOfData": "file",
    "CrawlerSchemaDeserializerVersion": "1.0",
    "CrawlerSchemaSerializerVersion": "1.0",
    "UPDATED_BY_CRAWLER": "crawl_date_table",
    "recordCount": "9"
},
"Compressed": false,
"BucketColumns": [],
"InputFormat": "org.apache.hadoop.mapred.TextInputFormat",
"SerdeInfo": {
    "Parameters": {
        "field.delim": ",",
    },
    "SerializationLibrary": "org.apache.hadoop.hive.serde2.lazy.LazySimpleSerDe"
}
}
}
]
}

```

2. Edit the table definition in your editor to add `"TableType": "EXTERNAL_TABLE"` to the table definition, as shown in the following example.

```

{
  "Table": {
    "Retention": 0,
    "TableType": "EXTERNAL_TABLE",
    "PartitionKeys": [

```

```

{
  "Name": "year",
  "Type": "string"
},
{
  "Name": "month",
  "Type": "string"
},
{
  "Name": "day",
  "Type": "string"
}
],
"UpdateTime": 1522368588.0,
"Name": "table_missing_table_type",
"StorageDescriptor": {
  "BucketColumns": [],
  "SortColumns": [],
  "StoredAsSubDirectories": false,
  "OutputFormat": "org.apache.hadoop.hive ql.io.HiveIgnoreKeyTextOutputFormat",
  "SerdeInfo": {
    "SerializationLibrary": "org.apache.hadoop.hive.serde2.lazy.LazySimpleSerDe",
    "Parameters": {
      "field.delim": ",",
    }
  },
},
"Parameters": {
  "classification": "csv",
  "CrawlerSchemaSerializerVersion": "1.0",
  "UPDATED_BY_CRAWLER": "crawl_date_table",
  "columnsOrdered": "true",
  "averageRecordSize": "7",
  "objectCount": "1",
  "sizeKey": "75",
  "delimiter": ",",
  "compressionType": "none",
  "recordCount": "9",
  "CrawlerSchemaDeserializerVersion": "1.0",
  "typeOfData": "file",
  "skip.header.line.count": "1"
},
"Columns": [
  {
    "Name": "col1",

```

```

    "Type": "string"
  },
  {
    "Name": "col2",
    "Type": "bigint"
  }
],
"Compressed": false,
"InputFormat": "org.apache.hadoop.mapred.TextInputFormat",
"NumberOfBuckets": -1,
"Location": "s3://myAthenatest/test_date_part/"
},
"Owner": "owner",
"Parameters": {
  "classification": "csv",
  "CrawlerSchemaSerializerVersion": "1.0",
  "UPDATED_BY_CRAWLER": "crawl_date_table",
  "columnsOrdered": "true",
  "averageRecordSize": "7",
  "objectCount": "1",
  "sizeKey": "75",
  "delimiter": ",",
  "compressionType": "none",
  "recordCount": "9",
  "CrawlerSchemaDeserializerVersion": "1.0",
  "typeOfData": "file",
  "skip.header.line.count": "1"
},
"LastAccessTime": 1513804142.0
}
}

```

3. You can adapt the following script to update the table input, so that it includes the `TableType` attribute.

```
aws glue update-table --database-name <your_datebase_name> --table-input
<updated_table_input>
```

The following shows an example.

```
aws glue update-table --database-name test_database --table-input '
{
  "Retention": 0,
```

```

"TableType": "EXTERNAL_TABLE",
"PartitionKeys": [
  {
    "Name": "year",
    "Type": "string"
  },
  {
    "Name": "month",
    "Type": "string"
  },
  {
    "Name": "day",
    "Type": "string"
  }
],
"Name": "table_missing_table_type",
"StorageDescriptor": {
  "BucketColumns": [],
  "SortColumns": [],
  "StoredAsSubDirectories": false,
  "OutputFormat": "org.apache.hadoop.hive ql.io.HiveIgnoreKeyTextOutputFormat",
  "SerdeInfo": {
    "SerializationLibrary": "org.apache.hadoop.hive.serde2.lazy.LazySimpleSerDe",
    "Parameters": {
      "field.delim": ",",
    }
  },
},
"Parameters": {
  "classification": "csv",
  "CrawlerSchemaSerializerVersion": "1.0",
  "UPDATED_BY_CRAWLER": "crawl_date_table",
  "columnsOrdered": "true",
  "averageRecordSize": "7",
  "objectCount": "1",
  "sizeKey": "75",
  "delimiter": ",",
  "compressionType": "none",
  "recordCount": "9",
  "CrawlerSchemaDeserializerVersion": "1.0",
  "typeOfData": "file",
  "skip.header.line.count": "1"
},
"Columns": [
  {

```

```
"Name": "col1",
  "Type": "string"
},
{
  "Name": "col2",
  "Type": "bigint"
}
],
"Compressed": false,
"InputFormat": "org.apache.hadoop.mapred.TextInputFormat",
"NumberOfBuckets": -1,
"Location": "s3://myAthenatest/test_date_part/"
},
"Owner": "owner",
"Parameters": {
  "classification": "csv",
  "CrawlerSchemaSerializerVersion": "1.0",
  "UPDATED_BY_CRAWLER": "crawl_date_table",
  "columnsOrdered": "true",
  "averageRecordSize": "7",
  "objectCount": "1",
  "sizeKey": "75",
  "delimiter": ",",
  "compressionType": "none",
  "recordCount": "9",
  "CrawlerSchemaDeserializerVersion": "1.0",
  "typeOfData": "file",
  "skip.header.line.count": "1"
},
"LastAccessTime": 1513804142.0
}'
```

Table not found when using Athena with Amazon Quick Sight

You can receive a "table not found" error if the tables in an analysis are missing from the Athena data source.

In the Athena console (<https://console.aws.amazon.com/athena/>), check for your table under the corresponding schema. You can recreate the table in Athena and then create a new dataset in Amazon Quick Sight on that table. To investigate how the table was lost in the first place, you can use the Athena console to check the query history. Doing this helps you find the queries that dropped the table.

If this error happened when you were editing a custom SQL query in preview, verify that the name of the table in the query, and check for any other syntax errors. Amazon Quick Sight can't infer the schema from the query. The schema must be specified in the query.

For example, the following statement works.

```
select from my_schema.my_table
```

The following statement fails because it's missing the schema.

```
select from my_table
```

If you still have the issue, verify that your tables, columns, and queries comply with Athena requirements. For more information, see [Names for Tables, Databases, and Columns](#) and [Troubleshooting](#) in the *Athena User Guide*.

Workgroup and output errors when using Athena with Quick Sight

To verify that workgroups are set up properly, check the following settings:

- **The Athena workgroup that's associated with the data source must exist.**

To fix this, you can return to the Athena data source settings and choose a different workgroup. For more information, see [Setting Up Workgroups](#) in the *Athena User Guide*.

Another solution is to have the AWS account administrator recreate the workgroup in the Athena console.

- **The Athena workgroup that's associated with the data source must be enabled.**

An AWS account administrator needs to enable the workgroup in the Athena console. Open the Athena console by using this direct link: <https://console.aws.amazon.com/athena/>. Then choose the appropriate workgroup in the **Workgroup** panel and view its settings. Choose **Enable workgroup**.

- **Make sure that you have access to the Amazon S3 output location that's associated with the Athena workgroup.**

To grant Amazon Quick Sight permissions to access the S3 output location, the Amazon Quick Sight administrator can edit **Security & Permissions** in the **Manage QuickSight** screen.

- **The Athena workgroup must have an associated S3 output location.**

An AWS account administrator needs to associate an S3 bucket with the workgroup in the Athena console. Open the Athena console by using this direct link: <https://console.aws.amazon.com/athena/>. Then choose the appropriate workgroup in the **Workgroup** panel and view its settings. Set **Query result location**.

Data source connectivity issues for Amazon Quick Sight

Use the following section to help you troubleshoot connections to data sources. Before you continue, verify that your database is currently available. Also, verify that you have the correct connection information and valid credentials.

Topics

- [I can't connect although my data source connection options look right \(SSL\)](#)
- [I can't connect to Amazon Athena](#)
- [I can't connect to Amazon S3](#)
- [I can't create or refresh a dataset from an existing Adobe Analytics data source](#)
- [I need to validate the connection to my data source, or change data source settings](#)
- [I can't connect to MySQL \(issues with SSL and authorization\)](#)
- [I can't connect to RDS](#)

I can't connect although my data source connection options look right (SSL)

Problems connecting can occur when Secure Sockets Layer (SSL) is incorrectly configured. The symptoms can include the following:

- You can connect to your database in other ways or from other locations but not in this case.
- You can connect to a similar database but not this one.

Before continuing, rule out the following circumstances:

- Permissions issues
- Availability issues
- An expired or invalid certificate

- A self-signed certificate
- Certificate chain in the wrong order
- Ports not enabled
- Firewall blocking an IP address
- Web Sockets are blocked
- A virtual private cloud (VPC) or security group not configured correctly.

To help find issues with SSL, you can use an online SSL checker, or a tool like OpenSSL.

The following steps walk through troubleshooting a connection where SSL is suspect. The administrator in this example has already installed OpenSSL.

Example

1. The user finds an issue connecting to the database. The user verifies that they can connect a different database in another AWS Region. They check other versions of the same database and can connect easily.
2. The administrator reviews the issue and decides to verify that the certificates are working correctly. The administrator searches online for an article on using OpenSSL to troubleshoot or debug SSL connections.
3. Using OpenSSL, the administrator verifies the SSL configuration in the terminal.

```
echo quit  
openssl s_client -connect <host>:port
```

The result shows that the certificate is not working.

```
...  
...  
...  
CONNECTED(00000003)  
012345678901234:error:140770FC:SSL routines:SSL23_GET_SERVER_HELLO:unknown  
  protocol:s23_clnt.c:782:  
---  
no peer certificate available
```

```
---
No client certificate CA names sent
---
SSL handshake has read 7 bytes and written 278 bytes
---
New, (NONE), Cipher is (NONE)
Secure Renegotiation IS NOT supported
SSL-Session:
    Protocol  : TLSv1.2
    Cipher    : 0000
    Session-ID:
    Session-ID-ctx:
    Master-Key:
    Key-Arg   : None
    PSK identity: None
    PSK identity hint: None
    Start Time: 1497569068
    Timeout   : 300 (sec)
    Verify return code: 0 (ok)
---
```

4. The administrator corrects the problem by installing the SSL certificate on the user's database server.

For more detail on the solution in this example, see [Using SSL to Encrypt a Connection to a DB Instance](#) in the *Amazon RDS User Guide*.

I can't connect to Amazon Athena

Intended audience: Amazon Quick administrators

Use this section to help troubleshoot connecting to Athena.

If you can't connect to Amazon Athena, you might get an insufficient permissions error when you run a query, showing that the permissions aren't configured. To verify that you can connect Amazon Quick Sight to Athena, check the following settings:

- AWS resource permissions inside of Amazon Quick Sight
- AWS Identity and Access Management (IAM) policies

- Amazon S3 location
- Query results location
- AWS KMS key policy (for encrypted datasets only)

For details, see following. For information about troubleshooting other Athena issues, see [Connectivity issues when using Amazon Athena with Amazon Quick Sight](#).

Make sure that you authorized Amazon Quick Sight to use Athena

Intended audience: Amazon Quick administrators

Use the following procedure to make sure that you successfully authorized Amazon Quick Sight to use Athena. Permissions to AWS resources apply to all Amazon Quick Sight users.

To perform this action, you must be an Amazon Quick Sight administrator. To check if you have access, verify that you see the **Manage QuickSight** option when you open the menu from your profile at upper right.

To authorize Amazon Quick Sight to access Athena

1. Choose your profile name (upper right). Choose **Manage Quick Sight**, and then scroll down to the **Custom permissions** section.
2. Choose **AWS resources** then choose **Add or remove**.
3. Find Athena in the list. Clear the box by Athena, then select it again to enable Athena.

Then choose **Connect both**.

4. Choose the buckets that you want to access from Amazon Quick Sight.

The settings for S3 buckets that you access here are the same ones that you access by choosing Amazon S3 from the list of AWS services. Be careful that you don't inadvertently disable a bucket that someone else uses.

5. Choose **Finish** to confirm your selection. Or choose **Cancel** to exit without saving.
6. Choose **Update** to save your new settings for Amazon Quick Sight access to AWS services. Or choose **Cancel** to exit without making any changes.
7. Make sure that you are using the correct AWS Region when you are finished.

If you had to change your AWS Region as part of the first step of this process, change it back to the AWS Region that you were using before.

Make sure that your IAM policies grant the right permissions

Intended audience: System administrators

Your AWS Identity and Access Management (IAM) policies must grant permissions to specific actions. Your IAM user or role must be able to read and write both the input and the output of the S3 buckets that Athena uses for your query.

If the dataset is encrypted, the IAM user needs to be a key user in the specified AWS KMS key's policy.

To verify that your IAM policies have permission to use S3 buckets for your query

1. Open the IAM console at <https://console.aws.amazon.com/iam/>.
2. Locate the IAM user or role that you are using. Choose the user or role name to see the associated policies.
3. Verify that your policy has the correct permissions. Choose a policy that you want to verify, and then choose **Edit policy**. Use the visual editor, which opens by default. If you have the JSON editor open instead, choose the **Visual editor** tab.
4. Choose the S3 entry in the list to see its contents. The policy needs to grant permissions to list, read, and write. If S3 is not in the list, or it doesn't have the correct permissions, you can add them here.

For examples of IAM policies that work with Quick Sight, see [IAM policy examples for Quick](#).

Make sure that the IAM user has read/write access to your S3 location

Intended audience: Amazon Quick administrators

To access Athena data from Quick Sight, first make sure that Athena and its S3 location are authorized in **Manage QuickSight** screen. For more information, see [Make sure that you authorized Amazon Quick Sight to use Athena](#).

Next, verify the relevant IAM permissions. The IAM user for your Athena connection needs read/write access to the location where your results go in S3. Start by verifying that the IAM user has an attached policy that [allows access to Athena](#), such as AmazonAthenaFullAccess. Let Athena create the bucket using the name that it requires, and then add this bucket to the list of buckets that QuickSight can access. If you change the default location of the results bucket (aws-athena-query-results-*), be sure that the IAM user has permission to read and write to the new location.

Verify that you don't include the AWS Region code in the S3 URL. For example, use s3://awsexamplebucket/path and not s3://us-east-1.amazonaws.com/awsexamplebucket/path. Using the wrong S3 URL causes an Access Denied error.

Also verify that the bucket policies and object access control lists (ACLs) [allow the IAM user to access the objects in the buckets](#). If the IAM user is in a different AWS account, see [Cross-account Access](#) in the *Amazon Athena User Guide*.

If the dataset is encrypted, verify that the IAM user is a key user in the specified AWS KMS key's policy. You can do this in the AWS KMS console at <https://console.aws.amazon.com/kms>.

To set permissions to your Athena query results location

1. Open the Athena console at <https://console.aws.amazon.com/athena/>.
2. Verify that you have selected the workgroup you want to use:
 - Examine the **Workgroup** option at the top. It has the format **Workgroup: *group-name***. If the group name is the one that you want to use, skip to the next step.
 - To choose a different workgroup, chose **Workgroup** at the top. Choose the workgroup that you want to use, and choose **Switch workgroup**.
3. Choose **Settings** at upper right.

(Not common) If you get an error that your workgroup is not found, use these steps to fix it:

- a. Ignore the error message for now, and instead find **Workgroup: *group-name*** on the **Settings** page. Your workgroup's name is a hyperlink. Open it.

- b. On the **Workgroup: <groupname>** page, choose **Edit workgroup** at left. Now close the error message.
- c. Near **Query result location**, open the S3 location selector by choosing the **Select** button that has the file folder icon.
- d. Choose the small arrow at the end of the name of the S3 location for Athena. The name must begin with `aws-athena-query-results`.
- e. (Optional) Encrypt query results by selecting the **Encrypt results stored in S3** check box.
- f. Choose **Save** to confirm your choices.
- g. If the error doesn't reappear, return to **Settings**.

Occasionally, the error might appear again. If so, take the following steps:

1. Choose the workgroup and then choose **View details**.
 2. (Optional) To preserve your settings, take notes or a screenshot of the workgroup configuration.
 3. Choose **Create workgroup**.
 4. Replace the workgroup with a new one. Configure the correct S3 location and encryption options. Note the S3 location because you need it later.
 5. Choose **Save** to proceed.
 6. When you no longer need the original workgroup, disable it. Make sure to carefully read the warning that appears, because it tells you what you lose if you choose to disable it.
4. If you didn't get this by troubleshooting in the previous step, choose **Settings** at upper right and get the S3 location value shown as **Query result location**.
 5. If **Encrypt query results** is enabled, check whether it uses SSE-KMS or CSE-KMS. Note the key.
 6. Open the S3 console at <https://console.aws.amazon.com/s3/>, open the correct bucket, and then choose the **Permissions** tab.
 7. Check that your IAM user has access by viewing **Bucket Policy**.

If you manage access with ACLs, make sure that the access control lists (ACLs) are set up by viewing **Access Control List**.

8. If your dataset is encrypted (**Encrypt query results** is selected in the workgroup settings), make sure that the IAM user or role is added as a key user in that AWS KMS key's policy. You can access AWS KMS settings at <https://console.aws.amazon.com/kms>.

To grant access to the S3 bucket used by Athena

1. Open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. Choose the S3 bucket used by Athena in the **Query result location**.
3. On the **Permissions** tab, verify the permissions.

For more information, see the AWS Support article [When I run an Athena query, I get an "Access Denied" error](#).

I can't connect to Amazon S3

To successfully connect to Amazon S3, make sure that you configure authentication and create a valid manifest file inside the bucket you are trying to access. Also, make sure that the file described by the manifest is available.

To verify authentication, make sure that you authorized Amazon Quick Sight to access the S3 account. It's not enough that you, the user, are authorized. Amazon Quick Sight must be authorized separately.

To authorize Amazon Quick Sight to access your Amazon S3 bucket

1. In the AWS Region list at upper right, choose the US East (N. Virginia) Region. You use this AWS Region temporarily while you edit your account permissions.
2. Inside Amazon Quick Sight, choose your profile name (upper right). Choose **Manage Quick Sight**, and then scroll down to the **Custom permissions** section.
3. Choose **AWS resources** then choose **Add or remove**.
4. Locate Amazon S3 in the list. Choose one of the following actions to open the screen where you can choose S3 buckets:
 - If the check box is clear, select the check box next to Amazon S3.
 - If the check box is selected, choose **Details**, and then choose **Select S3 buckets**.
5. Choose the buckets that you want to access from Amazon Quick Sight. Then choose **Select**.
6. Choose **Update**.
7. If you changed your AWS Region during the first step of this process, change it back to the AWS Region that you want to use.

We strongly recommend that you make sure that your manifest file is valid. If Amazon Quick Sight can't parse your file, it gives you an error message. That might be something like "We can't parse the manifest file as valid JSON" or "We can't connect to the S3 bucket."

To verify your manifest file

1. Open your manifest file. You can do this directly from the Amazon S3 console at <https://console.aws.amazon.com/s3/>. Go to your manifest file and choose **Open**.
2. Make sure that the URI or URLs provided inside the manifest file indicate the file or files that you want connect to.
3. Make sure that your manifest file is formed correctly, if you use a link to the manifest file rather than uploading the file. The link shouldn't have any additional phrases after the word `.json`. You can get the correct link to an S3 file by viewing its **Link** value in the details on the S3 console.
4. Make sure that the content of the manifest file is valid by using a JSON validator, like the one at <https://jsonlint.com>.
5. Verify permissions on your bucket or file. In the <https://console.aws.amazon.com/s3/>, navigate to your Amazon S3 bucket, choose the **Permissions** tab, and add the appropriate permissions. Make sure that the permissions are at the right level, either on the bucket or on the file or files.
6. If you are using the `s3://` protocol, rather than `https://`, make sure that you reference your bucket directly. For example, use `s3://awsexamplebucket/myfile.csv` instead of `s3://s3-us-west-2.amazonaws.com/awsexamplebucket/myfile.csv`. Doubly specifying Amazon S3, by using `s3://` and also `s3-us-west-2.amazonaws.com`, causes an error.

For more information about manifest files and connecting to Amazon S3, see [Supported formats for Amazon S3 manifest files](#).

In addition, verify that your Amazon S3 dataset was created according to the steps in [Creating a dataset using Amazon S3 files](#).

If you use Athena to connect to Amazon S3, see [I can't connect to Amazon Athena](#).

I can't create or refresh a dataset from an existing Adobe Analytics data source

As of May 1, 2022, Quick Sight no longer supports legacy OAuth and version 1.3 and SOAP API operations in Adobe Analytics. If you experience failures while trying to create or refresh a dataset from an existing Adobe Analytics data source, you might have a stale access token.

To troubleshoot failures while creating or refreshing a dataset from an existing Adobe Analytics data source

1. Open Quick Sight and choose **Data** at left.
2. Choose **New** then **Dataset**.
3. On the **Create a dataset** page, choose the Adobe Analytics data source that you want to update from the list of existing data sources.
4. Choose **Edit data source**.
5. On the **Edit Adobe Analytics data source** page that opens, choose **Update data source** to reauthorize the Adobe Analytics connection.
6. Try recreating or refreshing the dataset again. The dataset creation or refresh should succeed.

I need to validate the connection to my data source, or change data source settings

In some cases, you might need to update your data source, or you got a connection error and need to check your settings. If so, take the following steps.

To validate your connection to the data source

1. From the Quick Sight homepage, choose **Data** at left.
2. Choose **New** then **Dataset**.
3. You will see a list of existing data sources.
4. Choose the data source that you want to test or change.
5. If the option is offered, choose **Edit/Preview data**.
6. Choose **Validate connection**.
7. Make any changes that you want to make, then choose **Update data source**.

I can't connect to MySQL (issues with SSL and authorization)

To check on some common connection issues in MySQL, use the following steps. This procedure helps you find out if you have enabled SSL and granted usage rights.

To find solutions for some common connection issues in MySQL

1. Check `/etc/my.cnf` to make sure SSL is enabled for MySQL.

2. In MySQL, run the following command.

```
show status like 'Ssl%';
```

If SSL is working, you see results like the following.

Variable_name	Value
Ssl_accept_renegotiates	0
Ssl_accepts	1
Ssl_callback_cache_hits	0
Ssl_cipher	
Ssl_cipher_list	
Ssl_client_connects	0
Ssl_connect_renegotiates	0
Ssl_ctx_verify_depth	18446744073709551615
Ssl_ctx_verify_mode	5
Ssl_default_timeout	0
Ssl_finished_accepts	0
Ssl_finished_connects	0
Ssl_session_cache_hits	0
Ssl_session_cache_misses	0
Ssl_session_cache_mode	SERVER
Ssl_session_cache_overflows	0
Ssl_session_cache_size	128
Ssl_session_cache_timeouts	0
Ssl_sessions_reused	0
Ssl_used_session_cache_entries	0
Ssl_verify_depth	0
Ssl_verify_mode	0
Ssl_version	

If SSL is disabled, you see results like the following.

```
+-----+-----+
```

Variable_name	Value
Ssl_accept_renegotiates	0
Ssl_accepts	0
Ssl_callback_cache_hits	0
Ssl_cipher	
Ssl_cipher_list	
Ssl_client_connects	0
Ssl_connect_renegotiates	0
Ssl_ctx_verify_depth	0
Ssl_ctx_verify_mode	0
Ssl_default_timeout	0
Ssl_finished_accepts	0
Ssl_finished_connects	0
Ssl_session_cache_hits	0
Ssl_session_cache_misses	0
Ssl_session_cache_mode	NONE
Ssl_session_cache_overflows	0
Ssl_session_cache_size	0
Ssl_session_cache_timeouts	0
Ssl_sessions_reused	0
Ssl_used_session_cache_entries	0
Ssl_verify_depth	0
Ssl_verify_mode	0
Ssl_version	

3. Make sure that you have installed a supported SSL certificate on the database server.
4. Grant usage for the specific user to connect using SSL.

```
GRANT USAGE ON *.* TO 'encrypted_user'@'%' REQUIRE SSL;
```

Note

TLS 1.2 for MySQL connections requires MySQL version 5.7.28 or higher. If your MySQL server enforces TLS 1.2 only (for example, `tls_version = TLSv1.2`) and the server version is below 5.7.28, the SSL handshake fails with a Communications link failure

error. To resolve this, upgrade your MySQL or Aurora MySQL database to version 5.7.28 or higher.

For more detail on the solution in this example, see the following:

- [SSL Support for MySQL DB Instances](#) in the *Amazon RDS User Guide*.
- [Using SSL to Encrypt a Connection to a DB Instance](#) in the *Amazon RDS User Guide*.
- [MySQL documentation](#)

I can't connect to RDS

For details on troubleshooting connections to Amazon RDS, see [Creating a dataset from a database](#).

You can also refer to the Amazon RDS documentation on troubleshooting connections, [Cannot Connect to Amazon RDS DB Instance](#).

Login issues with Quick Sight

Use the following section to help you troubleshoot login and access issues with the Quick Sight console.

Topics

- [Insufficient permissions when using Athena with Amazon Quick Sight](#)
- [Amazon Quick Sight isn't working in my browser](#)
- [How do I delete my Amazon Quick Sight account?](#)
- [Individuals in my organization get an "External Login is Unauthorized" message when they try to access Quick Sight](#)
- [My email sign-in stopped working](#)

Insufficient permissions when using Athena with Amazon Quick Sight

If you receive an error message that says you have insufficient permissions, try the following steps to resolve your problem.

You need administrator permissions to troubleshoot this issue.

To resolve an insufficient permissions error

1. Make sure that Amazon Quick Sight can access the Amazon S3 buckets used by Athena:
 - a. To do this, choose your profile name (upper right). Choose **Manage Quick Sight**, and then scroll down to the **Custom permissions** section.
 - b. Choose **AWS resources** then choose **Add or remove**.
 - c. Locate Athena in the list. Clear the check box by Athena, then select it again to enable Athena.

Choose **Connect both**.

- d. Choose the buckets that you want to access from Amazon Quick Sight.

The settings for S3 buckets that you access here are the same ones that you access by choosing Amazon S3 from the list of AWS services. Be careful that you don't inadvertently disable a bucket that someone else uses.

- e. Choose **Select** to save your S3 buckets.
 - f. Choose **Update** to save your new settings for Amazon Quick Sight access to AWS services. Or choose **Cancel** to exit without making any changes.
2. If your data file is encrypted with an AWS KMS key, grant permissions to the Amazon Quick Sight IAM role to decrypt the key. The easiest way to do this is to use the AWS CLI.

You can run the [create-grant](#) command in AWS CLI to do this.

```
aws kms create-grant --key-id <AWS KMS key ARN> --grantee-principal <Your Amazon Quick Sight Role ARN> --operations Decrypt
```

The Amazon Resource Name (ARN) for the Amazon Quick Sight role has the format `arn:aws:iam::<account id>:role/service-role/aws-quicksight-service-role-v<version number>` and can be accessed from the IAM console. To find your AWS KMS key ARN, use the S3 console. Go to the bucket that contains your data file and choose the **Overview** tab. The key is located near **KMS key ID**.

For Amazon Athena, Amazon S3, and Athena Query Federation connections, Quick Sight uses the following IAM role by default:

```
arn:aws:iam::AWS-ACCOUNT-ID:role/service-role/aws-quicksight-s3-consumers-role-v0
```

If the `aws-quicksight-s3-consumers-role-v0` is not present, then Quick Sight uses:

```
arn:aws:iam::AWS-ACCOUNT-ID:role/service-role/aws-quicksight-service-role-v0
```

Amazon Quick Sight isn't working in my browser

If you can't view Amazon Quick Sight correctly in your Google Chrome browser, take the following steps to fix the problem.

To view Amazon Quick Sight in your Chrome browser

1. Open Chrome and go to `chrome://flags/#touch-events`.
2. If the option is set to **Automatic**, change it to **Disabled**.
3. Close and reopen Chrome.

How do I delete my Amazon Quick Sight account?

In some cases, you might need to delete your Amazon Quick Sight account even when you can't access Amazon Quick Sight to unsubscribe. If so, sign in to AWS and use the following link to open [the unsubscribe screen](https://us-east-1.quicksight.aws.amazon.com/sn/console/unsubscribe): `https://us-east-1.quicksight.aws.amazon.com/sn/console/unsubscribe`. This approach works no matter what AWS Regions that you use. It deletes all data, analyses, Amazon Quick Sight users, and Amazon Quick Sight administrators. If you have further difficulty, contact support.

Individuals in my organization get an "External Login is Unauthorized" message when they try to access Quick Sight

Intended audience: Amazon Quick administrators

When an individual in your organization is federating into Quick Sight using **AssumeRoleWithWebIdentity**, Quick Sight maps a single role-based user to a single external login. In some cases, that individual might be authenticated through an external login (such as Amazon Cognito) that's different from the originally mapped user. If so, they can't access Quick Sight and get the following unexpected error message.

The external login used for federation is unauthorized for the Quick Sight user.

To learn how to troubleshoot this issue, see the following sections:

- [Why is this happening?](#)
- [How can I fix it?](#)

Why is this happening?

You are using a simplified Amazon Cognito flow

If you're using Amazon Cognito to federate into Quick Sight, the single sign-on (IAM Identity Center) setup might use the `CognitoIdentityCredentials` API operation to assume the Quick Sight role. This method maps all users in the Amazon Cognito identity pool to a single Quick Sight user and isn't supported by Quick Sight.

We recommend that you use the `AssumeRoleWithWebIdentity` API operation instead, which specifies the role session name.

You're using unauthenticated Amazon Cognito users

Amazon Cognito IAM Identity Center is set up for unauthenticated users in the Amazon Cognito identity pool. The Quick Sight role trust policy is set up like the following example.

This setup allows a temporary Amazon Cognito user to assume a role session mapped to a unique Quick Sight user. Because unauthenticated identities are temporary, they aren't supported by Quick Sight.

We recommend that you don't use this setup, which setup isn't supported by Quick Sight. For Quick Sight, make sure that the Amazon Cognito IAM Identity Center uses authenticated users.

You deleted and recreated an Amazon Cognito user with the same user name attributes

In this case, the associated Amazon Cognito user that's mapped to the Quick Sight user was deleted and recreated. The newly created Amazon Cognito user has a different underlying subject. Depending on how the role session name is mapped to the Quick Sight user, the session name might correspond to the same Quick Sight role-based user.

We recommend that you remap the Quick Sight user to the updated Amazon Cognito user subject by using the `UpdateUser` API operation. For more information, see the following [UpdateUser API example](#).

You're mapping multiple Amazon Cognito user pools in different AWS accounts to one identity pool and with Quick Sight

Mapping multiple Amazon Cognito user pools in different AWS accounts to one identity pool and Quick Sight isn't supported by Quick Sight.

How can I fix it?

You can use Quick Sight public API operations to update the external login information for your users. Use the following options to learn how.

Use RegisterUser to create users with external login information

If the external login provider is Amazon Cognito, use the following CLI code to create users.

```
aws quicksight register-user --aws-account-id account-id --namespace namespace --email user-email --user-role user-role --identity-type IAM --iam-arn arn:aws:iam::account-id:role/cognito-associated-iam-role --session-name cognito-username --external-login-federation-provider-type COGNITO --external-login-id cognito-identity-id --region identity-region
```

The external-login-id should be the identity ID for the Amazon Cognito user. The format is <identity-region>:<cognito-user-sub>, as shown in the following example.

```
aws quicksight register-user --aws-account-id 111222333 --namespace default --email cognito-user@amazon.com --user-role ADMIN --identity-type IAM --iam-arn arn:aws:iam::111222333:role/CognitoQuickSightRole --session-name cognito-user --external-login-federation-provider-type COGNITO --external-login-id us-east-1:12345678-1234-1234-abc1-a1b1234567 --region us-east-1
```

If the external login provider is a custom OpenID Connect (OIDC) provider, use the following CLI code to create users.

```
aws quicksight register-user --aws-account-id account-id --namespace namespace --email user-email --user-role user-role --identity-type IAM --iam-arn arn:aws:iam::account-id:role/identity-provider-associated-iam-role --session-name identity-username --external-login-federation-provider-type CUSTOM_OIDC --custom-federation-provider-url custom-identity-provider-url --external-login-id custom-provider-identity-id --region identity-region
```

The following is an example.


```
aws quicksight register-user --aws-account-id 111222333 --namespace default
--email identity-user@amazon.com --user-role ADMIN --identity-type IAM
--iam-arn arn:aws:iam::111222333:role/CustomIdentityQuickSightRole
--session-name identity-user --external-login-federation-provider-type CUSTOM_OIDC
--custom-federation-provider-url idp.us-east-1.amazonaws.com/us-east-1_ABCDE
--external-login-id 12345678-1234-1234-abc1-a1b1234567 --region us-east-1
```

To learn more about using RegisterUser in the CLI, see [RegisterUser](#) in the *Amazon Quick API Reference*.

Use DescribeUser to check external login information for users

If a user is a role-based federated user from an external login provider, use the DescribeUser API operation to check the external login information for it, as shown in the following code.

```
aws quicksight describe-user --aws-account-id account-id --namespace namespace
--user-name identity-provider-associated-iam-role/identity-username
--region identity-region
```

The following is an example.

```
aws quicksight describe-user --aws-account-id 111222333 --namespace default --user-name
IdentityQuickSightRole/user --region us-west-2
```

The result contains the external login information fields if there are any. Following is an example.

```
{
  "Status": 200,
  "User": {
    "Arn": "arn:aws:quicksight:us-east-1:111222333:user-default-
IdentityQuickSightRole-user",
    "UserName": "IdentityQuickSightRole-user",
    "Email": "user@amazon.com",
    "Role": "ADMIN",
    "IdentityType": "IAM",
    "Active": true,
    "PrincipalId": "federated-iam-AR0AAAAAAAAAAAAAAAAA:user",
    "ExternalLoginFederationProviderType": "COGNITO",
    "ExternalLoginFederationProviderUrl": "cognito-identity.amazonaws.com",
    "ExternalLoginId": "us-east-1:123abc-1234-123a-b123-12345678a"
  },
}
```

```
"RequestId": "12345678-1234-1234-abc1-a1b1234567"
}
```

To learn more about using `DescribeUser` in the CLI, see [DescribeUser](#) in the *Amazon Quick API Reference*.

Use `UpdateUser` to update external login information for users

In some cases, you might find that the external login information saved for the user from the `DescribeUser` result isn't correct or the external login information is missing. If so, you can use the `UpdateUser` API operation to update it. Use the following examples.

For Amazon Cognito users, use the following.

```
aws quicksight update-user --aws-account-id account-id --namespace namespace
--user-name cognito-associated-iam-role/cognito-username
--email user-email --role user-role
--external-login-federation-provider-type COGNITO
--external-login-id cognito-identity-id --region identity-region
```

The following is an example.

```
aws quicksight update-user --aws-account-id 111222333 --namespace default
--user-name CognitoQuickSightRole/cognito-user --email cognito-user@amazon.com
--role ADMIN --external-login-federation-provider-type COGNITO
--external-login-id us-east-1:12345678-1234-1234-abc1-a1b1234567 --region us-west-2
```

For custom OIDC provider users, use the following.

```
aws quicksight update-user --aws-account-id account-id --namespace namespace
--user-name identity-provider-associated-iam-role/identity-username
--email user-email --role user-role
--external-login-federation-provider-type CUSTOM_OIDC
--custom-federation-provider-url custom-identity-provider-url
--external-login-id custom-provider-identity-id --region identity-region
```

The following is an example.

```
aws quicksight update-user --aws-account-id 111222333 --namespace default
--user-name IdentityQuickSightRole/user --email user@amazon.com --role ADMIN
```

```
--external-login-federation-provider-type CUSTOM_OIDC
--custom-federation-provider-url idp.us-east-1.amazonaws.com/us-east-1_ABCDE
--external-login-id 123abc-1234-123a-b123-12345678a --region us-west-2
```

If you want to delete the external login information for the user, use `NONE` external login federation provider type. Use the following CLI command to delete external login information.

```
aws quicksight update-user --aws-account-id account-id --namespace namespace
--user-name identity-provider-associated-iam-role/identity-username
--email user-email --role user-role
--external-login-federation-provider-type NONE --region identity-region
```

The following is an example.

```
aws quicksight update-user --aws-account-id 111222333 --namespace default
--user-name CognitoQuickSightRole/cognito-user --email cognito-user@amazon.com --role
ADMIN --external-login-federation-provider-type NONE --region us-west-2
```

To learn more about using `UpdateUser` in the CLI, see the [UpdateUser](#) in the *Amazon Quick API Reference*.

My email sign-in stopped working

Currently, emails are case-sensitive. If yours isn't working, ask your administrator to check it for a mix of upper and lowercase letters. Use your email as it was entered.

Visual issues with Quick Sight

Use the following section to help you troubleshoot problems with visuals and their formatting.

Topics

- [I can't see my visuals](#)
- [I get a feedback bar across my printed documents](#)
- [My map charts don't show locations](#)
- [My pivot table stops working](#)
- [My visual can't find missing columns](#)

- [My visual can't find the query table](#)
- [My visual doesn't update after I change a calculated field](#)
- [Values in a Microsoft Excel file with scientific notation don't format correctly in Quick Sight](#)

I can't see my visuals


Use the following section to help you troubleshoot missing visuals. Before you continue, check to make sure you can still access your data source. If you can't connect to your data source, see [Data source connectivity issues for Amazon Quick Sight](#).

- If you are having trouble adding a visual to an analysis, try the following:
 - Check your connectivity and confirm that you have access to all domains that Quick Sight uses for access. To see a list of all URLs Quick Sight uses, see [Domains accessed by Quick Sight](#).
 - Check that you aren't trying to add more objects than the quota allows. Amazon Quick Sight supports up to 30 datasets in a single analysis, up to 30 visuals in a single sheet, and a limit of 20 sheets per analysis.
 - Suppose that you are editing an analysis for a selected data source and the connection to the data source ends unexpectedly. The resulting error state can prevent further changes to the analysis. In this case, you can't add more visuals to the analysis. Check for this state.
- If your visuals don't load, try the following:
 - If you are using a corporate network, seek out help from your network administrator and verify that the network's firewall settings permit traffic from *.aws.amazon.com, amazonaws.com, wss://*.aws.amazon.com, and cloudfront.net.
 - Add exceptions to your ad blocker for *.aws.amazon.com, amazonaws.com, wss://*.aws.amazon.com, and cloudfront.net.
 - If you are using a proxy server, verify that *.quicksight.aws.amazon.com and cloudfront.net are added to the list of approved domains (the allow list).

I get a feedback bar across my printed documents

The browser sometimes prints the document feedback bar across the page, blocking some printed content.

To avoid this problem, use the twirl-down icon on the bottom left of the screen (shown following) to minimize the feedback bar. Then print your document.



My map charts don't show locations

For automatic mapping, called geocoding, to work on map charts, make sure that your data is prepared following specific rules. For help with geospatial issues, see [Geospatial troubleshooting](#). For help with preparing data for geospatial charts, see [Adding geospatial data](#).

My pivot table stops working

If your pivot table exceeds the computational limitations of the underlying database, this is usually caused by the combination of items in the field wells. That is, it's caused by a combination of rows, columns, metrics, and table calculations. To reduce the level of complexity and the potential for errors, simplify your pivot table. For more information, see [Pivot table best practices](#).

My visual can't find missing columns

The visuals in my analysis aren't working as expected. The error message says "The column(s) used in this visual do not exist."

The most common cause of this error is that your data source schema changed. For example, it's possible a column name changed from a_column to b_column.

Depending on how your dataset accesses the data source, choose one of the following.

- If the dataset is based on custom SQL, do one or more of the following:
 - Edit the dataset.
 - Edit the SQL statement.

For example, if the table name changed from a_column to b_column, you can update the SQL statement to create an alias: `SELECT b_column as a_column`. By using the alias to maintain the same field name in the dataset, you avoid having to add the column to your visuals as a new entity.

When you're done, choose **Save & visualize**.

- If the dataset isn't based on custom SQL, do one or more of the following:
 - Edit the dataset.

- For fields that now have different names, rename them in the dataset. You can use the field names from your original dataset. Then open your analysis and add the renamed fields to the affected visuals.

When you're done, choose **Save & visualize**.

My visual can't find the query table

In this case, the visuals in your analysis aren't working as expected. The error message says "Amazon Quick Sight can't find the query table."

The most common cause of this error is that your data source schema changed. For example, it's possible a table name changed from `x_table` to `y_table`.

Depending on how the dataset accesses the data source, choose one of the following.

- If the dataset is based on custom SQL, do one or more of the following:
 - Edit the dataset.
 - Edit the SQL statement.

For example, if the table name changed from `x_table` to `y_table`, you can update the FROM clause in the SQL statement to refer to the new table instead.

When you're done, choose **Save & visualize**, then choose each visual and readd the fields as needed.

- If the dataset isn't based on custom SQL, do the following:
 1. Create a new dataset using the new table, `y_table` for example.
 2. Open your analysis.
 3. Replace the original dataset with the newly created dataset. If there are no column changes, all the visuals should work after you replace the dataset. For more information, see [Replacing datasets](#).

My visual doesn't update after I change a calculated field

When you update a calculated field that many other fields depend on, the consuming entities might not update as expected. For example, when you update a calculated field that's used by a field being visualized, the visual doesn't update as expected.

To resolve this issue, refresh your internet browser.

Values in a Microsoft Excel file with scientific notation don't format correctly in Quick Sight

When you connect to a Microsoft Excel file that has a number column that contains values with scientific notation, they might not format correctly in Quick Sight. For example, the value 1.59964E+11, which is actually 159964032802, formats as 159964000000 in Quick Sight. This can lead to an incorrect analysis.

To resolve this issue, format the column as Text in Microsoft Excel, and then upload the file to Quick Sight.

Developing with Amazon Quick Sight

We provide API operations for Amazon Quick Sight, and also software development kits (SDKs) for AWS that enable you to access Amazon Quick Sight from your preferred programming language. Currently, you can manage users and groups. In Enterprise edition, you can also embed dashboards in your webpage or app.

To monitor the calls made to the Amazon Quick Sight API for your account, including calls made by the AWS Management Console, command line tools, and other services, use AWS CloudTrail. For more information, see the [AWS CloudTrail User Guide](#).

Required knowledge

If you plan to access Amazon Quick Sight through an API, you should be familiar with the following:

- JSON
- Web services
- HTTP requests
- One or more programming languages, such as JavaScript, Java, Python, or C#.

We recommend visiting the AWS [Getting Started Resource Center](#) for a tour of what AWS SDKs and toolkits have to offer.

Although you can use a terminal and your favorite text editor, you might benefit from the more visual UI experience you get in an integrated development environment (IDE). We provide a list

of IDEs in the *AWS Getting Started Resource Center* in the [IDE and IDE Toolkits](#) section. This site provides AWS toolkits that you can download for your preferred IDE. Some IDEs also offer tutorials to help you learn more about programming languages.

Available API operations for Amazon Quick Sight

AWS provides libraries, sample code, tutorials, and other resources for software developers who prefer to build applications using language-specific API operations instead of submitting a request over HTTPS. These libraries provide basic functions that automatically take care of tasks such as cryptographically signing your requests, retrying requests, and handling error responses. These libraries help make it easier for you to get started.

For more information about downloading the AWS SDKs, see [AWS SDKs and Tools](#). The following links are a sample of the language-specific API documentation available.

AWS Command Line Interface

- [AWS CLI QuickSight Command Reference](#)
- [AWS CLI User Guide](#)
- [AWS CLI Command Reference](#)

AWS SDK for .NET

- [Amazon.Quicksight](#)
- [Amazon.Quicksight.Model](#)

AWS SDK for C++

- [Aws::QuickSight::QuickSightClient Class Reference](#)

AWS SDK for Go

- [quicksight](#)

AWS SDK for Java

- [com.amazonaws.services.quicksight](#)
- [com.amazonaws.services.quicksight.model](#)

AWS SDK for JavaScript

- [AWS.QuickSight](#)

AWS SDK for PHP

- [QuickSightClient](#)

AWS SDK for Python (Boto3)

- [Amazon Quick Sight](#)

AWS SDK for Ruby

- [Aws::QuickSight](#)

Terminology and concepts

This section provides a list of terms for development in Amazon Quick Sight.

Anonymous Amazon Quick Sight User: – A temporary Amazon Quick Sight user identity that virtually belongs to a namespace, and is usable only with embedding. You can use tag-based rules to implement row-level security for such users.

Caller identity: – The identity of the AWS Identity and Access Management user making an API request. The identity of the caller is determined by Amazon Quick Sight using the signature attached to the request. Through the use of our provided SDK clients, no manual steps are necessary to generate the signature or attach it to the requests. However, you can do it manually if you want to.

Invoker identity: – In addition to the caller identity, but not as a replacement for it, you can assume a caller's identity through the IAM AssumeRole API when making calls to Amazon Quick Sight. AWS approves callers through their invoker's identity. This is done to avoid having to explicitly add multiple accounts belonging to the same Amazon Quick Sight subscription.

Namespace: – a logical container that allows you to isolate user pools so that you can organize clients, subsidiaries, teams, and so on. For more information, see [Supporting multitenancy with isolated namespaces](#)

QuickSight ARN: – Amazon Resource Name (ARN). Amazon Quick Sight resources are identified using their name or ARN. For example, these are the ARNs for a group named MyGroup1, a user named User1, and a dashboard with the ID 1a1ac2b2-3fc3-4b44-5e5d-c6db6778df89:

```
arn:aws:quicksight:us-east-1:111122223333:group/default/MyGroup1
arn:aws:quicksight:us-east-1:111122223333:user/default/User1
arn:aws:quicksight:us-west-2:111122223333:dashboard/1a1ac2b2-3fc3-4b44-5e5d-
c6db6778df89
```

The following examples show ARNs for a template named MyTemplate and a dashboard named MyDashboard.

1. Sample ARN for a template

```
arn:aws:quicksight:us-east-1:111122223333:template/MyTemplate
```

2. Sample ARN for a template, referencing a specific version of the template

```
arn:aws:quicksight:us-east-1:111122223333:template/MyTemplate/version/10
```

3. Sample ARN for a template alias

```
arn:aws:quicksight:us-east-1:111122223333:template/MyTemplate/alias/STAGING
```

4. Sample ARN for a dashboard

```
arn:aws:quicksight:us-east-1:111122223333:dashboard/MyDashboard
```

5. Sample ARN for a dashboard, referencing a specific version of the dashboard

```
arn:aws:quicksight:us-east-1:111122223333:dashboard/MyDashboard/version/10
```

Depending on the scenario, you might need to provide an entity's name, ID, or ARN. You can retrieve the ARN if you have the name, using some of the Amazon Quick Sight API operations.

Amazon Quick Sight dashboard: – An entity which identifies Amazon Quick Sight reports, created from analyses or templates. Amazon Quick Sight dashboards are sharable. With the right permissions, scheduled email reports can be created from them. The CreateDashboard and DescribeDashboard API Operations act on the dashboard entity.

Amazon Quick Sight template: – An entity which encapsulates the metadata required to create an analysis or a dashboard. It abstracts the dataset associated with the analysis by replacing it with placeholders. Templates can be used to create dashboards by replacing dataset placeholders with datasets that follow the same schema that was used to create the source analysis and template.

Amazon Quick Sight user: – This is an Amazon Quick Sight user identity acted upon by your API call. This user isn't identical to the caller identity but might be the one that maps to the user within Amazon Quick Sight.

Developing applications with the Amazon Quick Sight API

You can manage most aspects of your deployment by using the AWS SDKs to access an API that's tailored to the programming language or platform that you're using. For more information, see [AWS SDKs](#).

For more information on the API operations, see [Amazon Quick Sight API Reference](#).

Before you can call the Amazon Quick Sight API operations, you need the `quicksight:operation-name` permission in a policy attached to your IAM identity. For example, to call `list-users`, you need the permission `quicksight:ListUsers`. The same pattern applies to all operations.

If you're not sure what the necessary permission is, you can attempt to make a call. The client then tells you what the missing permission is. You can use asterisk (*) in the Resource field of your permission policy instead of specifying explicit resources. However, we recommended that you restrict each permission as much as possible. You can restrict user access by specifying or excluding resources in the policy, using their Amazon Quick Sight Amazon Resource Name (ARN) identifier.

For more information, see the following:

- [IAM policy examples for Amazon Quick Sight](#)
- [Actions, Resources, and Condition Keys](#)
- [IAM JSON Policy Elements](#)

To retrieve the ARN of a user or a group, use the `Describe` operation on the relevant resource. You can also add conditions in IAM to further restrict access to an API in some scenarios. For instance, when adding `User1` to `Group1`, the main resource is `Group1`, so you can allow or deny access to certain groups, but you can also add a condition by using the IAM Amazon Quick Sight key `quicksight:UserName` to allow or prevent certain users from being added to that group.

Following is an example policy. It means that the caller with this policy attached, is able to invoke the `CreateGroupMembership` operation on any group, provided that the user name they are adding to the group is not `user1`.

```
{
  "Effect": "Allow",
  "Action": "quicksight:CreateGroupMembership",
  "Resource": "arn:aws:quicksight:us-east-1:aws-account-id:group/default/*",
  "Condition": {
    "StringNotEquals": {
      "quicksight:UserName": "user1"
    }
  }
}
```

AWS CLI

The following procedure explains how to interact with Amazon Quick Sight API operations through the AWS CLI. The following instructions have been tested in Bash but should be identical or similar in other command-line environments.

1. Install AWS SDK in your environment. Instructions on how to do that are located here: [AWS Command line Interface](#).
2. Set up your AWS CLI identity and region using the following command and follow-up instructions. Use the credentials for an IAM identity or role that has the proper permissions.

```
aws configure
```

3. Look at the Amazon Quick Sight SDK help by issuing the following command:

```
aws quicksight help
```

4. To get detailed instructions on how to use an API, enter its name followed by help, like so:

```
aws quicksight list-users help
```

5. Now you can call an Amazon Quick Sight API operation. This example returns a list of Amazon Quick Sight users in your account.

```
aws quicksight list-users --aws-account-id aws-account-id --namespace default --  
region us-east-1
```

Java SDK

Use the following procedure to set up a Java app that interacts with Amazon Quick Sight.

1. To get started, create a Java project in your IDE.
2. Import the Amazon Quick Sight SDK into your new project, for example:
`AWSQuickSightJavaClient-1.11.x.jar`
3. Once your IDE indexes the Amazon Quick Sight SDK, you should be able to add an import line as follows:

```
import com.amazonaws.services.quicksight.AmazonQuickSight;
```

If your IDE doesn't recognize this as valid, verify that you imported the SDK.

4. Like other AWS SDKs, Amazon Quick Sight SDK requires external dependencies to perform many of its functions. You need to download and import those into the same project. The following dependencies are required:
 - `aws-java-sdk-1.11.402.jar` (AWS Java SDK and credentials setup) — See [Set up the AWS SDK for Java](#)
 - `commons-logging-1.2.jar` — See https://commons.apache.org/proper/commons-logging/download_logging.cgi
 - `jackson-annotations-2.9.6.jar`, `jackson-core-2.9.6.jar`, and `jackson-databind-2.9.6.jar` — See <http://repo1.maven.org/maven2/com/fasterxml/jackson/core/>
 - `httpclient-4.5.6.jar`, `httpcore-4.4.10.jar` — See <https://hc.apache.org/downloads.cgi>
 - `joda-time-2.1.jar` — See <https://mvnrepository.com/artifact/joda-time/joda-time/2.1>
5. Now, you are ready to create an Amazon Quick Sight client. You can use a default public endpoint that the client can communicate with or you can reference the endpoint explicitly. There are multiple ways to provide your AWS credentials. In the following example, we

provide a direct, simple approach. The following client method is used to make all the API calls that follow:

```
private static AmazonQuickSight getClient() {
    final AWSCredentialsProvider credsProvider = new AWSCredentialsProvider() {
        @Override
        public AWSCredentials getCredentials() {
            // provide actual IAM access key and secret key here
            return new BasicAWSCredentials("access-key", "secret-key");
        }

        @Override
        public void refresh() {}
    };

    return AmazonQuickSightClientBuilder
        .standard()
        .withRegion(Regions.US_EAST_1.getName())
        .withCredentials(credsProvider)
        .build();
}
```

6. Now, we can use the above client to list all the users in our Amazon Quick Sight account.

 Note

You have to provide the AWS account ID that you used to subscribe to Amazon Quick Sight. This must match the AWS account ID of the caller's identity. Cross-account calls aren't supported at this time. Furthermore, the required parameter namespace should always be set to *default*.

```
getClient().listUsers(new ListUsersRequest()
    .withAwsAccountId("relevant_AWS_account_ID")
    .withNamespace("default"))
    .getUserList().forEach(user -> {
        System.out.println(user.getArn());
    });
```

7. To see a list of all possible API operations and the request objects they use, you can **CTRL-click** on the client object in your IDE in order to view the Amazon Quick Sight interface.

Alternatively, find it within the `com.amazonaws.services.quicksight` package in the Amazon Quick Sight JavaClient JAR file.

JavaScript (Node.js) SDK

Use the following procedure to interact with Amazon Quick Sight using Node.js.

1. Set up your node environment using the following commands:
 - `npm install aws-sdk`
 - `npm install aws4`
 - `npm install request`
 - `npm install url`
2. For information on configuring the Node.js with AWS SDK and setting your credentials, see--> the [AWS SDK for JavaScript Developer Guide for SDK v2](#).
3. Use the following code sample to test your setup. HTTPS is required. The sample displays a full listing of Amazon Quick Sight operations along with their URL request parameters, followed by a list of Amazon Quick Sight users in your account.

```
const AWS = require('aws-sdk');
const https = require('https');

var quicksight = new AWS.Service({
  apiConfig: require('./quicksight-2018-04-01.min.json'),
  region: 'us-east-1',
});

console.log(quicksight.config.apiConfig.operations);

quicksight.listUsers({
  // Enter your actual AWS account ID
  'AwsAccountId': 'relevant_AWS_account_ID',
  'Namespace': 'default',
}, function(err, data) {
  console.log('---');
  console.log('Errors: ');
  console.log(err);
  console.log('---');
  console.log('Response: ');
  console.log(data);
});
```

```
});
```

Python3 SDK

Use the following procedure to create a custom built botocore package to interact with Amazon Quick Sight.

1. Create a credentials file in the AWS directory for your environment. In a Linux/Mac-based environment, that file is called `~/.aws/credentials` and looks like this:

```
[default]
aws_access_key_id = Your_IAM_access_key
aws_secret_access_key = Your_IAM_secret_key
```

2. Unzip the folder `botocore-1.12.10`. Change directory into `botocore-1.12.10` and enter the Python3 interpreter environment.
3. Responses come back as a dictionary object. They each have a `ResponseMetadata` entry that contains request IDs and response status. Other entries are based on what type of operation you run.
4. The following example is a sample app that first creates, deletes, and lists groups. Then, it lists users in a Quicksight account:

```
import botocore.session
default_namespace = 'default'
account_id = 'relevant_AWS_Account'

session = botocore.session.get_session()
client = session.create_client("quicksight", region_name='us-east-1')

print('Creating three groups: ')
client.create_group(AwsAccountId = account_id, Namespace=default_namespace,
    GroupName='MyGroup1')
client.create_group(AwsAccountId = account_id, Namespace=default_namespace,
    GroupName='MyGroup2')
client.create_group(AwsAccountId = account_id, Namespace=default_namespace,
    GroupName='MyGroup3')

print('Retrieving the groups and listing them: ')
response = client.list_groups(AwsAccountId = account_id,
    Namespace=default_namespace)
```



```
for group in response['GroupList']:
    print(group)

print('Deleting our groups: ')
client.delete_group(AwsAccountId = account_id, Namespace=default_namespace,
    GroupName='MyGroup1')
client.delete_group(AwsAccountId = account_id, Namespace=default_namespace,
    GroupName='MyGroup2')
client.delete_group(AwsAccountId = account_id, Namespace=default_namespace,
    GroupName='MyGroup3')

response = client.list_users(AwsAccountId = account_id,
    Namespace=default_namespace)
for user in response['UserList']:
    print(user)
```

.NET/C# SDK

Use the following procedure to interact with Amazon Quick Sight using C#.NET. This example is constructed on Microsoft Visual for Mac; the instructions can vary slightly based on your IDE and platform. However, they should be similar.

1. Unzip the `nuget.zip` file into a folder called `nuget`.
2. Create a new **Console app** project in Visual Studio.
3. Under your solution, locate app **Dependencies**, then open the context (right-click menu) and choose **Add Packages**.
4. In the sources list, choose **Configure Sources**.
5. Choose **Add**, and name the source `QuickSightSDK`. Browse to the `nuget` folder and choose **Add Source**.
6. Choose **OK**. Then, with `QuickSightSDK` selected, select all three Amazon Quick Sight packages:
 - `AWSSDK.QuickSight`
 - `AWSSDK.Extensions.NETCore.Setup`
 - `AWSSDK.Extensions.CognitoAuthentication`
7. Click **Add Package**.
8. Copy and paste the following sample app into your console app editor.

```
using System;
using Amazon.QuickSight.Model;
using Amazon.QuickSight;

namespace DotNetQuickSightSDKTest
{
    class Program
    {
        private static readonly string AccessKey = "insert_your_access_key";
        private static readonly string SecretAccessKey =
            "insert_your_secret_key";
        private static readonly string AccountID = "AWS_account_ID";
        private static readonly string Namespace = "default"; // leave this as
        default

        static void Main(string[] args)
        {
            var client = new AmazonQuickSightClient(
                AccessKey,
                SecretAccessKey,
                Amazon.RegionEndpoint.USEast1);

            var listUsersRequest = new ListUsersRequest
            {
                AwsAccountId = AccountID,
                Namespace = Namespace
            };

            client.ListUsersAsync(listUsersRequest).Result.UserList.ForEach(
                user => Console.WriteLine(user.Arn)
            );

            var listGroupsRequest = new ListGroupsRequest
            {
                AwsAccountId = AccountID,
                Namespace = Namespace
            };

            client.ListGroupsAsync(listGroupsRequest).Result.GroupList.ForEach(
                group => Console.WriteLine(group.Arn)
            );
        }
    }
}
```

```
}

```

Amazon Quick Sight events integration

With Amazon EventBridge, you can respond automatically to events in Amazon Quick Sight such as new dashboard creation or updates. These events are delivered to EventBridge in near real time. As a developer, you can write simple rules to indicate which events are of interest, and what actions to take when an event matches a rule. By using events, you can complete use cases such as continuous backup and deployment.

Topics

- [Supported events](#)
- [Example event payload](#)
- [Creating rules to send Amazon Quick Sight events to Amazon CloudWatch](#)
- [Creating rules to send Amazon Quick Sight events to AWS Lambda](#)

Supported events

Amazon Quick Sight currently supports the following events.

Asset type	Action	Event detail type	Event detail
Dashboard	Create	Amazon Quick Sight Dashboard Creation Successful	<pre>{ "dashboardId": "6fdbbc328-ebbd-457f-aa02-9780173afc83", "versionNumber": 1 }</pre>
Dashboard	Create	Amazon Quick Sight Dashboard Creation Failed	<pre>{ "dashboardId": "6fdbbc328-ebbd-457f-</pre>

Asset type	Action	Event detail type	Event detail
			<pre>aa02-9780173afc83", "versionNumber": 1, "errors": [{ "Type": "PARAMETER_NOT_FOUND", "Message": "Missing property abc" }, { "Type": "DATA_SET_NOT_FOUND", "Message": "Cannot find dataset with id abc" }] }</pre>
Dashboard	Create	Amazon Quick Sight Dashboard Permissions Updated	<pre>{"dashboardId": "6fdb328-ebbd-457f-aa02-9780173afc83" }</pre>

Asset type	Action	Event detail type	Event detail
Dashboard	Update	Amazon Quick Sight Dashboard Update Successful	<pre>{ "dashboardId": "6fdb328-ebbd-457f-aa02-9780173afc83", "versionNumber": 1 }</pre>
Dashboard	Update	Amazon Quick Sight Dashboard Update Failed	<pre>{ "dashboardId": "6fdb328-ebbd-457f-aa02-9780173afc83", "versionNumber": 1, "errors": [{ "Type": "PARAMETER_NOT_FOUND", "Message": "Missing property abc" }, { "Type": "DATA_SET_NOT_FOUND", "Message": "Cannot find dataset with id abc" }] }</pre>

Asset type	Action	Event detail type	Event detail
Dashboard	Update	Amazon Quick Sight Dashboard Permissions Updated	<pre>{ "dashboardId": "6fdbbc328- ebbd-457f- aa02-9780173afc8 3"} }</pre>
Dashboard	Publish	Amazon Quick Sight Dashboard Published Version Updated	<pre>{ "dashboardId": "6fdbbc328- ebbd-457f- aa02-9780173afc8 3", "versionNumber": 2 }</pre>
Dashboard	Delete	Amazon Quick Sight Dashboard Deleted	<pre>{ "dashboardId": "6fdbbc328- ebbd-457f- aa02-9780173afc8 3" }</pre>
Analysis	Create	Amazon Quick Sight Analysis Creation Successful	<pre>{ "analysisId": "e5f37119- e24c-4874-901a- af9032b729b5" }</pre>

Asset type	Action	Event detail type	Event detail
Analysis	Create	Amazon Quick Sight Analysis Creation Failed	<pre>{ "analysisId": "e5f37119-e24c-4874-901a-af9032b729b5", "errors": [{ "Type": "PARAMETER_NOT_FOUND", "Message": "Missing property abc" }, { "Type": "DATA_SET_NOT_FOUND", "Message": "Cannot find dataset with id abc" }] }</pre>
Analysis	Create	Amazon Quick Sight Analysis Permissions Updated	<pre>{"analysisId": "e5f37119-e24c-4874-901a-af9032b729b5" }</pre>
Analysis	Delete	Amazon Quick Sight Analysis Deleted	<pre>{ "analysisId": "e5f37119-e24c-4874-901a-af9032b729b5" }</pre>

Asset type	Action	Event detail type	Event detail
Analysis	Update	Amazon Quick Sight Analysis update successful	<pre>{ "analysis Id": "e5f37119- e24c-4874-901a- af9032b729b5" }</pre>

Asset type	Action	Event detail type	Event detail
Analysis	Update	Amazon Quick Sight Analysis update failed	<pre>{ "analysis Id": "e5f37119- e24c-4874-901a- af9032b729b5", "errors": [{ "Type": "PARAMETE R_NOT_FOUND", "Message" : "Missing property abc" }, { "Type": "DATA_SET _NOT_FOUND", "Message": "Cannot find dataset with id abc" }] }</pre>

Asset type	Action	Event detail type	Event detail
Analysis	Update	Amazon Quick Sight Analysis Permissions Updated	<pre>{ "analysisId": "e5f37119- e24c-4874-901a- af9032b729b5" }</pre>
VPC connection	Create	Amazon Quick Sight VPC Connection Creation Successful	<pre>{ "vpcConne ctionId": "53d34238 -57e7-488d- b99a-a0037d275a4 e", "availabi lityStatu s": "CREATION _SUCCESSFUL" }</pre>
VPC connection	Create	Amazon Quick Sight VPC Connection Creation Failed	<pre>{ "vpcConne ctionId": "53d34238 -57e7-488d- b99a-a0037d275a4 e", "availabi lityStatu s": "CREATION _FAILED" }</pre>

Asset type	Action	Event detail type	Event detail
VPC connection	Update	Amazon Quick Sight VPC Connection Update Successful	<pre>{ "vpcConnectionId": "53d34238-57e7-488d-b99a-a0037d275a4e", "availabilityStatus": "UPDATE_SUCCESSFUL" }</pre>
VPC connection	Update	Amazon Quick Sight VPC Connection Update Failed	<pre>{ "vpcConnectionId": "53d34238-57e7-488d-b99a-a0037d275a4e", "availabilityStatus": "UPDATE_FAILED" }</pre>
VPC connection	Delete	Amazon Quick Sight VPC Connection Deletion Successful	<pre>{ "vpcConnectionId": "53d34238-57e7-488d-b99a-a0037d275a4e", "availabilityStatus": "DELETED" }</pre>

Asset type	Action	Event detail type	Event detail
VPC connection	Delete	Amazon Quick Sight VPC Connection Deletion Failed	<pre>{ "vpcConnectionId": "53d34238-57e7-488d-b99a-a0037d275a4e", "availabilityStatus": "DELETION_FAILED" }</pre>
Folder	Create	Amazon Quick Sight Folder Created	<pre>{ "folderId": "77e307e8-b41b-472a-90e8-fe3f471537be", "parentFolderArn": "arn:aws:quicksight:us-east-1:123456789012:folder/098765432134" }</pre>
Folder	Create	Amazon Quick Sight Folder Permissions Updated	<pre>{"folderId": "77e307e8-b41b-472a-90e8-fe3f471537be" }</pre>

Asset type	Action	Event detail type	Event detail
Folder	Update	Amazon Quick Sight Folder Updated	<pre>{ "folderId": "77e307e8-b41b-472a-90e8-fe3f471537be" }</pre>
Folder	Update	Amazon Quick Sight Folder Permissions Updated	<pre>{"folderId": "77e307e8-b41b-472a-90e8-fe3f471537be" }</pre>
Folder	Delete	Amazon Quick Sight Folder Deleted	<pre>{ "folderId": "77e307e8-b41b-472a-90e8-fe3f471537be" }</pre>
Folder	Membership update	Amazon Quick Sight Folder Membership Updated	<pre>{ "folderId": "77e307e8-b41b-472a-90e8-fe3f471537be", "membersAdded": ["arn:aws:quicksight:us-east-1:123456789012:analysis/5f37119-e24c-4874-901a-af9032b729b5"], "membersRemoved": [] }</pre>

Asset type	Action	Event detail type	Event detail
Dataset	Create	Amazon Quick Sight Dataset Created	<pre>{ "datasetId": "a6553a81-f97e-4ffa-a860-baea63196efa" }</pre>
Dataset	Create	Amazon Quick Sight Dataset Permissions Updated	<pre>{"datasetId": "a6553a81-f97e-4ffa-a860-baea63196efa" }</pre>
Dataset	Update	Amazon Quick Sight Dataset Updated	<pre>{ "datasetId": "a6553a81-f97e-4ffa-a860-baea63196efa" }</pre>
Dataset	Update	Amazon Quick Sight Dataset Permissions Updated	<pre>{"datasetId": "a6553a81-f97e-4ffa-a860-baea63196efa" }</pre>
Dataset	Delete	Amazon Quick Sight Dataset Deleted	<pre>{ "datasetId": "a6553a81-f97e-4ffa-a860-baea63196efa" }</pre>

Asset type	Action	Event detail type	Event detail
DataSource	Create	Amazon Quick Sight DataSource Creation Successful	<pre>{ "datasourceId": "230caa6e-dc87-406b-91fb-037f29c32824" }</pre>
DataSource	Create	Amazon Quick Sight DataSource Creation Failed	<pre>{ "datasourceId": "230caa6e-dc87-406b-91fb-037f29c32824", "error": { "message": "AMAZON_ELASTICSEARCH engine version 7.4 is lower than minimum supported version 7.7", "type": "ENGINE_VERSION_NOT_SUPPORTED" } }</pre>
DataSource	Create	Amazon Quick Sight DataSource Permissions Updated	<pre>{"datasourceId": "230caa6e-dc87-406b-91fb-037f29c32824" }</pre>

Asset type	Action	Event detail type	Event detail
DataSource	Update	Amazon Quick Sight DataSource Update Successful	<pre>{ "datasourceId": "230caa6e-dc87-406b-91fb-037f29c32824" }</pre>
DataSource	Update	Amazon Quick Sight DataSource Update Failed	<pre>{ "datasourceId": "230caa6e-dc87-406b-91fb-037f29c32824", "error": { "message": "AMAZON_ELASTICSEARCH engine version 7.4 is lower than minimum supported version 7.7", "type": "ENGINE_VERSION_NOT_SUPPORTED" } }</pre>
DataSource	Update	Amazon Quick Sight DataSource Permissions Updated	<pre>{"datasourceId": "230caa6e-dc87-406b-91fb-037f29c32824" }</pre>

Asset type	Action	Event detail type	Event detail
DataSource	Delete	Amazon Quick Sight DataSource Deleted	<pre>{ "datasourceId": "230caa6e-dc87-406b-91fb-037f29c32824" }</pre>
Theme	Create	Amazon Quick Sight Theme Creation Successful	<pre>{ "themeId": "6fdb328-ebbd-457f-aa02-9780173afc83", "versionNumber": 1 }</pre>
Theme	Create	Amazon Quick Sight Theme Creation Failed	<pre>{ "themeId": "6fdb328-ebbd-457f-aa02-9780173afc83", "versionNumber": 1 }</pre>
Theme	Create	Amazon Quick Sight Theme Permissions Updated	<pre>{"themeId": "6fdb328-ebbd-457f-aa02-9780173afc83" }</pre>

Asset type	Action	Event detail type	Event detail
Theme	Update	Amazon Quick Sight Theme Update Successful	<pre>{ "themeId" : "6fdbbc328-ebbd-457f-aa02-9780173afc83", "versionNumber": 2 }</pre>
Theme	Update	Amazon Quick Sight Theme Update Failed	<pre>{ "themeId" : "6fdbbc328-ebbd-457f-aa02-9780173afc83", "versionNumber": 2 }</pre>
Theme	Update	Amazon Quick Sight Theme Permissions Updated	<pre>{"themeId": "6fdbbc328-ebbd-457f-aa02-9780173afc83" }</pre>
Theme	Delete	Amazon Quick Sight Theme deleted	<pre>{ "themeId" : "6fdbbc328-ebbd-457f-aa02-9780173afc83" }</pre>

Asset type	Action	Event detail type	Event detail
Theme	Alias Create	Amazon Quick Sight Theme Alias Created	<pre>{ "themeId" : "6fdb328 -ebbd-457f- aa02-9780173afc8 3", "aliasName": "MyThemeAlias" "versionN umber": 2 }</pre>
Theme	Alias Update	Amazon Quick Sight Alias Updated	<pre>{ "themeId" : "6fdb328 -ebbd-457f- aa02-9780173afc8 3", "aliasName": "MyThemeAlias" "versionN umber": 4 }</pre>
Theme	Alias Delete	Amazon Quick Sight Theme Alias Deleted	<pre>{ "themeId" : "6fdb328 -ebbd-457f- aa02-9780173afc8 3", "aliasName": "MyThemeAlias" "versionN umber": 2 }</pre>

Example event payload

All events follow the standard EventBridge [object structure](#). The detail field is a JSON object that contains more information about the event.

```
{
  "version": "0",
  "id": "3acb26c8-397c-4c89-a80a-ce672a864c55",
  "detail-type": "QuickSight Dashboard Creation Successful",
  "source": "aws.quicksight",
  "account": "123456789012",
  "time": "2023-10-30T22:06:31Z",
  "region": "us-east-1",
  "resources": ["arn:aws:quicksight:us-east-1:123456789012:dashboard/6fdb328-ebbd-457f-aa02-9780173afc83"],
  "detail": {
    "dashboardId": "6fdb328-ebbd-457f-aa02-9780173afc83",
    "versionNumber": 1
  }
}
```

Creating rules to send Amazon Quick Sight events to Amazon CloudWatch

You can write simple rules to indicate which Amazon Quick Sight events interest you and which automated actions to take when an event matches a rule. For example, you can configure Amazon Quick Sight to send events to Amazon CloudWatch whenever a Amazon Quick Sight asset is placed in a folder. For more information, see the [Amazon EventBridge user guide](#).

1. Sign in to the AWS Management Console and open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. Under **Events** in the navigation pane, choose **Rules**.
3. Choose **Create rule**.
4. Enter a name and description for the rule. The rule name must be unique within this Region. For example, enter QuickSightAssetChangeRuleCloudWatch.
5. Choose **default** Event bus.
6. Choose **Rule with an event pattern**, and then choose **Next**.
7. For **Event source**, choose **AWS events or EventBridge partner events**.
8. In the **Creation method** section, choose **Custom pattern (JSON editor)**.

9. In the **Event pattern** text box, enter the following snippet and choose **Next**.

```
{
  "source": ["aws.quicksight"]
}
```

Alternatively, you can create the rule that only subscribes to a subset of event types in Amazon Quick Sight. For example, the following rule will only triggered when an asset is added to or removed from a folder with id 77e307e8-b41b-472a-90e8-fe3f471537be.

```
{
  "source": ["aws.quicksight"],
  "detail-type": ["QuickSight Folder Membership Updated"],
  "detail": {
    "folderId": "77e307e8-b41b-472a-90e8-fe3f471537be"
  }
}
```

10. For **Targets**, choose **AWS service > CloudWatch log group**.
11. Choose from an existing log group or create a new one by entering a new log group name.
12. Optionally, you can add another target for this rule.
13. In **Configure tags**, choose **Next**.
14. Choose **Create rule**.

For more information, see [Creating Amazon EventBridge rule that reacts To events](#) in the Amazon EventBridge user guide.

Creating rules to send Amazon Quick Sight events to AWS Lambda

In this tutorial, you create an AWS Lambda function that logs the asset events in the Amazon Quick Sight account. You then create a rule that runs the function whenever there is an asset change. This tutorial assumes that you already signed up for Amazon Quick Sight.

Step 1: Create a Lambda function

Create a Lambda function to log the state change events. You specify this function when you create your rule.

1. Sign in to the AWS Management Console and open the AWS Lambda console at <https://console.aws.amazon.com/lambda/>.
2. If you're new to Lambda, you see a welcome page. Choose **Get Started Now**. Otherwise, choose **Create function**.
3. Choose **Author from scratch**.
4. On the Create function page, enter a name and description for the Lambda function. For example, name the function `QuickSightAssetChangeFn`.
5. In **Runtime**, select **Node.js 18.x**.
6. For **Architecture**, choose **x86_64**.
7. For **Execution role**, choose either **Create a new role with basic Lambda permissions** or **Use an existing role** and choose the role you want.
8. Choose **Create function**.
9. On the **QuickSightAssetChange** page, choose **index.js**.
10. In the **index.js** pane, delete the existing code.
11. Enter the following code snippet.

```
console.log('Loading function');
exports.handler = async (event, context) => {
  console.log('Received QuickSight event:', JSON.stringify(event));
};
```

12. Choose **Deploy**.

Step 2: Create a rule

Create a rule to run your Lambda function whenever you create/update/delete a Amazon Quick Sight asset.

1. Sign in to the AWS Management Console and open the Amazon EventBridge console at <https://console.aws.amazon.com/events/>.
2. In the navigation pane, choose **Rules**.
3. Choose **Create rule**.
4. Enter a name and description for the rule. For example, enter `QuickSightAssetChangeRule`.
5. Select **default** Event bus.

6. Choose **Rule with an event pattern**, and then choose **Next**.
7. For **Event source**, choose **AWS events or EventBridge partner events**.
8. In the **Creation method** section, choose **Custom pattern (JSON editor)**.
9. In the **Event pattern** text box, enter the following snippet and choose **Next**.

```
{
  "source": ["aws.quicksight"]
}
```

Alternatively, you can create the rule that only subscribes to a subset of event types in Amazon Quick Sight. For example, the following rule will only triggered when an asset is added to or removed from a folder with id 77e307e8-b41b-472a-90e8-fe3f471537be.

```
{
  "source": ["aws.quicksight"],
  "detail-type": ["QuickSight Folder Membership Updated"],
  "detail": {
    "folderId": "77e307e8-b41b-472a-90e8-fe3f471537be"
  }
}
```

10. For **Target types**, choose **AWS service** and **Lambda function**.
11. For **Function**, choose the Lambda function that you created. Then choose **Next**.
12. In **Configure tags**, choose **Next**.
13. Review the steps in your rule. Then choose **Create rule**.

Step 3: Test the rule

To test your rule, create an analysis. After waiting a minute, verify that your Lambda function was invoked.

1. Open the Amazon Quick Sight console at <https://quicksight.aws.amazon.com/>.
2. Create a new analysis.
3. In the navigation pane, choose **Rules**, choose the name of the rule that you created.
4. In **Rule details**, choose **Monitoring**.
5. You will be redirected to the Amazon CloudWatch console. If you are not redirected, choose **View the metrics in CloudWatch**.

6. In **All metrics**, choose the name of the rule that you created. The graph indicates that the rule was invoked.
7. In the navigation pane, choose **Log groups**.
8. Choose the name of the log group for your Lambda function. For example, `/aws/lambda/function-name`.
9. Choose the name of the log stream to view the data provided by the function for the instance that you launched. You should see a received event similar to the following:

```
{
  "version": "0",
  "id": "3acb26c8-397c-4c89-a80a-ce672a864c55",
  "detail-type": "QuickSight Analysis Creation Successful",
  "source": "aws.quicksight",
  "account": "123456789012",
  "time": "2023-10-30T22:06:31Z",
  "region": "us-east-1",
  "resources": ["arn:aws:quicksight:us-east-1:123456789012:analysis/e5f37119-e24c-4874-901a-af9032b729b5"],
  "detail": {
    "analysisId": "e5f37119-e24c-4874-901a-af9032b729b5"
  }
}
```

For an example of Amazon Quick Sight event in JSON format, see [Overview of events for Amazon Quick Sight](#).

Embedded analytics for Amazon Quick Sight

Important

Amazon Quick Sight has new API operations for embedding analytics: `GenerateEmbedUrlForAnonymousUser` and `GenerateEmbedUrlForRegisteredUser`. You can still use the `GetDashboardEmbedUrl` and `GetSessionEmbedUrl` API operations to embed dashboards and the Amazon Quick Sight console, but they don't contain the latest embedding capabilities. For more information about embedding using the old API operations, see [Embedding analytics using the `GetDashboardEmbedURL` and `GetSessionEmbedURL` API operations](#).

Applies to: Enterprise Edition

Intended audience: Amazon Quick developers

With Amazon Quick Sight embedded analytics, you can seamlessly integrate data-driven experiences into your software applications. You can style the embedded components to match your brand. This capability brings the power of Amazon Quick Sight to your end users, where they can analyze and interact with data without ever leaving the application. Improving the user experience by reducing cognitive complexity gives users a better opportunity for deeper understanding and effectiveness.

Amazon Quick Sight supports embedding for these elements:

- Amazon Quick Sight console (full authoring experience for registered users)
- Amazon Quick Sight dashboards and visuals (for registered users, anonymous users, public end users)
- Amazon Quick Sight Q search bar (for registered users and anonymous users)

With an embedded Amazon Quick Sight console, you embed the full Amazon Quick Sight experience. Doing this makes it possible to use Amazon Quick Sight authoring tools as part of your application, rather than in the context of the AWS Management Console or a standalone website. Users of an embedded Amazon Quick Sight console need to be registered as Amazon Quick Sight authors or admins in your AWS account. They also need to be authenticated into the same AWS account, using any of the Amazon Quick Sight-supported authentication methods.

With an embedded Amazon Quick Sight dashboard or visual, readers get the same functionality and interactivity as they do in a published dashboard or visual. To use an embedded dashboard or visual, readers (viewers) can include any of the following:

- Amazon Quick Sight users authenticated in your AWS account by any method supported by Amazon Quick Sight.
- Unauthenticated visitors to a website or application – This option requires session packs with capacity pricing . For information about subscription types, see [Understanding Amazon Quick Sight subscriptions and roles](#).

- Multiple end users viewing a display on monitors or large screens by programmatic access.

If your app also resides in AWS, the app doesn't need to reside on the same AWS account as the Amazon Quick Sight subscription. However, the app needs to be able to assume the AWS Identity and Access Management (IAM) role that you use for the API calls.

Before you can embed content, make sure that you're using Amazon Quick Sight Enterprise edition in the AWS account where you plan to use embedding.

Amazon Quick Sight embedding is available in all supported AWS Regions.

Topics

- [Embedding Amazon Quick Sight analytics into your applications](#)
- [Embedding custom Amazon Quick Sight assets into your application](#)
- [Embedding Amazon Quick Sight visuals and dashboards with a 1-click embed code](#)
- [Embedding with the Amazon Quick Sight APIs](#)

Embedding Amazon Quick Sight analytics into your applications

Applies to: Enterprise Edition

To embed analytics, you can run the Amazon Quick Sight embedding API to generate the embed code. Alternatively for dashboards, you can copy an embed code when you share the dashboard in Amazon Quick Sight. Each option is described below.

1-click embedding for registered users

When you share a dashboard with registered users in your account, you can copy an embed code for the dashboard and paste it into your internal application's HTML.

Using 1-click enterprise embedding is best for when you want to embed a Amazon Quick Sight dashboard in an internal application that users need to authenticate in to. When you copy the embed code, you get a static embed code that doesn't change.

For more information, see [Embedding Amazon Quick Sight visuals and dashboards for registered users with a 1-click embed code](#).

Embedding with the Amazon Quick Sight APIs

Embedding with the Amazon Quick Sight API is best for when you want to embed the Amazon Quick Sight experience in an internal application that users must authenticate in to, or an external application that anyone can access. When you use the embedding API operations to generate an embed code, you get a one-time code.

For more information, see [Embedding with the Amazon Quick Sight APIs](#).

Embedding custom Amazon Quick Sight assets into your application

You can use Amazon Quick Sight embedded analytics to embed custom Amazon Quick Sight assets into your application that are tailored to meet your business needs. For embedded dashboards and visuals, Amazon Quick Sight authors can add filters and drill downs that readers can access as they navigate the dashboard or visual. Amazon Quick Sight developers can also use the Amazon Quick Sight SDKs to build tighter integrations between their SaaS applications and their Amazon Quick Sight embedded assets to add datapoint callback actions to visuals on a dashboard at runtime.

For more information about the Amazon Quick Sight SDKs, see the `amazon-quicksight-embedding-sdk` on [GitHub](#) or [NPM](#).

Following, you can find descriptions of how to use the Amazon Quick Sight SDKs to customize your Amazon Quick Sight embedded analytics.

Topics

- [Adding embedded callback actions at runtime in Amazon Quick Sight](#)
- [Filtering data at runtime for Amazon Quick Sight embedded dashboards and visuals](#)
- [Customize the look and feel of Amazon Quick Sight embedded dashboards and visuals](#)
- [Using the Amazon Quick Sight Embedding SDK to enable shareable links to embedded dashboard views](#)

Adding embedded callback actions at runtime in Amazon Quick Sight

Use embedded datapoint callback actions to build tighter integrations between your software as a service (SaaS) application and your Amazon Quick Sight embedded dashboards and visuals. Developers can register datapoints to be called back with the Amazon Quick Sight embedding SDK. When you register a callback action for a visual, readers can select a datapoint on the visual to receive a callback that provides data specific to the selected data point. This information can be

used to flag key records, compile raw data specific to the datapoint, capture records, and compile data for backend processes.

Embedded callbacks aren't supported for custom visual content, text boxes, or insights.

Before you begin registering datapoints for callback, update the Embedding SDK to version 2.3.0. For more information about using the Amazon Quick Sight Embedding SDK, see the [amazon-quicksight-embedding-sdk](#) on GitHub.

A datapoint callback can be registered to one or more visuals at runtime through the Amazon Quick Sight SDK. You can also register a datapoint callback to any interaction supported by the [VisualCustomAction](#) API structure. This allows the datapoint callback to initiate when the user selects the datapoint on the visual or when the datapoint is selected from the datapoint context menu. The following example registers a datapoint callback that the reader initiates when they select a datapoint on the visual.

```
/const MY_GET_EMBED_URL_ENDPOINT =  
  "https://my.api.endpoint.domain/MyGetEmbedUrlApi"; // Sample URL  
  
// The dashboard id to embed  
const MY_DASHBOARD_ID = "my-dashboard"; // Sample ID  
  
// The container element in your page that will have the embedded dashboard  
const MY_DASHBOARD_CONTAINER = "#experience-container"; // Sample ID  
  
// SOME HELPERS  
  
const ActionTrigger = {  
  DATA_POINT_CLICK: "DATA_POINT_CLICK",  
  DATA_POINT_MENU: "DATA_POINT_MENU",  
};  
  
const ActionStatus = {  
  ENABLED: "ENABLED",  
  DISABLED: "DISABLED",  
};  
  
// This function makes a request to your endpoint to obtain an embed url for a given  
// dashboard id  
// The example implementation below assumes the endpoint takes dashboardId as request  
// data  
// and returns an object with EmbedUrl property
```

```
const myGetEmbedUrl = async (dashboardId) => {
  const apiOptions = {
    dashboardId,
  };
  const apiUrl = new URL(MY_GET_EMBED_URL_ENDPOINT);
  apiUrl.search = new URLSearchParams(apiOptions).toString();
  const apiResponse = await fetch(apiUrl.toString());
  const apiResponseData = await apiResponse.json();
  return apiResponseData.EmbedUrl;
};

// This function constructs a custom action object
const myConstructCustomActionModel = (
  customActionId,
  actionName,
  actionTrigger,
  actionStatus
) => {
  return {
    Name: actionName,
    CustomActionId: customActionId,
    Status: actionStatus,
    Trigger: actionTrigger,
    ActionOperations: [
      {
        CallbackOperation: {
          EmbeddingMessage: {},
        },
      },
    ],
  };
};

// This function adds a custom action on the first visual of first sheet of the
// embedded dashboard
const myAddVisualActionOnFirstVisualOfFirstSheet = async (
  embeddedDashboard
) => {
  // 1. List the sheets on the dashboard
  const { SheetId } = (await embeddedDashboard.getSheets())[0];
  // If you'd like to add action on the current sheet instead, you can use
  // getSelectedSheetId method
  // const SheetId = await embeddedDashboard.getSelectedSheetId();
```

```

// 2. List the visuals on the specified sheet
const { VisualId } = (await embeddedDashboard.getSheetVisuals(SheetId))[0];

// 3. Add the custom action to the visual
try {
  const customActionId = "custom_action_id"; // Sample ID
  const actionName = "Flag record"; // Sample name
  const actionTrigger = ActionTrigger.DATA_POINT_CLICK; // or
ActionTrigger.DATA_POINT_MENU
  const actionStatus = ActionStatus.ENABLED;
  const myCustomAction = myConstructCustomActionModel(
    customActionId,
    actionName,
    actionTrigger,
    actionStatus
  );
  const response = await embeddedDashboard.addVisualActions(
    SheetId,
    VisualId,
    [myCustomAction]
  );
  if (!response.success) {
    console.log("Adding visual action failed", response.errorCode);
  }
} catch (error) {
  console.log("Adding visual action failed", error);
}
};

const parseDatapoint = (visualId, datapoint) => {
  datapoint.Columns.forEach((Column, index) => {
    // FIELD | METRIC
    const columnType = Object.keys(Column)[0];

    // STRING | DATE | INTEGER | DECIMAL
    const valueType = Object.keys(Column[columnType])[0];
    const { Column: columnMetadata } = Column[columnType][valueType];

    const value = datapoint.RawValues[index][valueType];
    const formattedValue = datapoint.FormattedValues[index];

    console.log(
      `Column: ${columnMetadata.ColumnName} has a raw value of ${value}
        and formatted value of ${formattedValue.Value} for visual: ${visualId}`
    );
  });
};

```

```

    );
  });
};

// This function is used to start a custom workflow after the end user selects a
// datapoint
const myCustomDatapointCallbackWorkflow = (callbackData) => {
  const { VisualId, Datapoints } = callbackData;

  parseDatapoint(VisualId, Datapoints);
};

// EMBEDDING THE DASHBOARD

const main = async () => {
  // 1. Get embed url
  let url;
  try {
    url = await myGetEmbedUrl(MY_DASHBOARD_ID);
  } catch (error) {
    console.log("Obtaining an embed url failed");
  }

  if (!url) {
    return;
  }

  // 2. Create embedding context
  const embeddingContext = await createEmbeddingContext();

  // 3. Embed the dashboard
  const embeddedDashboard = await embeddingContext.embedDashboard(
    {
      url,
      container: MY_DASHBOARD_CONTAINER,
      width: "1200px",
      height: "300px",
      resizeHeightOnSizeChangedEvent: true,
    },
    {
      onMessage: async (messageEvent) => {
        const { eventName, message } = messageEvent;
        switch (eventName) {
          case "CONTENT_LOADED": {

```

```

        await myAddVisualActionOnFirstVisualOfFirstSheet(embeddedDashboard);
        break;
    }
    case "CALLBACK_OPERATION_INVOKED": {
        myCustomDatapointCallbackWorkflow(message);
        break;
    }
    },
    },
    );
};

main().catch(console.error);

```

You can also configure the preceding example to initiate datapoint callback when the user opens the context menu. To do this with the preceding example, set the value of `actionTrigger` to `ActionTrigger.DATA_POINT_MENU`.

After a datapoint callback is registered, it's applied to most datapoints on the specified visual or visuals. Callbacks don't apply to totals or subtotals on visuals. When a reader interacts with a datapoint, a `CALLBACK_OPERATION_INVOKED` message is emitted to the Amazon Quick Sight embedding SDK. This message is captured by the `onMessage` handler. The message contains the raw and display values for the full row of data associated with the datapoint that is selected. It also contains the column metadata for all columns in the visual that the datapoint is contained in. The following is an example of a `CALLBACK_OPERATION_INVOKED` message.

```

{
  CustomActionId: "custom_action_id",
  DashboardId: "dashboard_id",
  SheetId: "sheet_id",
  VisualId: "visual_id",
  DataPoints: [
    {
      RawValues: [
        {
          String: "Texas" // 1st raw value in row
        },
        {
          Integer: 1000 // 2nd raw value in row
        }
      ],
    },
  ],
}

```



```

FormattedValues: [
    {Value: "Texas"}, // 1st formatted value in row
    {Value: "1,000"} // 2nd formatted value in row
],
Columns: [
    { // 1st column metadata
      Dimension: {
        String: {
          Column: {
            ColumnName: "State",
            DatsetIdentifier: "..."
          }
        }
      },
    },
    { // 2nd column metadata
      Measure: {
        Integer: {
          Column: {
            ColumnName: "Cancelled",
            DatsetIdentifier: "..."
          },
          AggregationFunction: {
            SimpleNumericalAggregation: "SUM"
          }
        }
      },
    },
  ],
]
}

```

Filtering data at runtime for Amazon Quick Sight embedded dashboards and visuals

You can use filter methods in the Amazon Quick Sight embedding SDK to leverage the power of Amazon Quick Sight filters within your software as a service (SaaS) application at runtime. Runtime filters allow business owners to integrate their application with their embedded Amazon Quick Sight dashboards and visuals. To accomplish this, create customized filter controls in your application and apply filter presets based on data from your application. Then, developers can personalize filter configurations for end users at runtime.

Developers can create, query, update, and remove Amazon Quick Sight filters on an embedded dashboard or visual from their application with the Amazon Quick Sight Embedding SDK. Create Amazon Quick Sight filter objects in your application with the [FilterGroup](#) data model and apply them to embedded dashboards and visuals using the filter methods. For more information about using the Amazon Quick Sight Embedding SDK, see the [amazon-quicksight-embedding-sdk](#) on GitHub.

Prerequisites

Before you can get started, make sure that you are using the Amazon Quick Sight Embedding SDK version 2.5.0 or higher.

Terminology and concepts

The following terminology can be useful when working with embedded runtime filtering.

- *Filter group* – A group of individual filters. Filters that are located within a `FilterGroup` are OR-ed with each other. Filters within a [FilterGroup](#) are applied to the same sheets or visuals.
- *Filter* – A single filter. The filter can be a category, numeric, or datetime filter type. For more information on filters, see [Filter](#).

Setting up

Before you begin, make sure that you have the following assets and information prepared.

- The sheet ID of the sheet that you want to scope the `FilterGroup` to. This can be obtained with the `getSheets` method in the Embedding SDK.
- The dataset and column identifier of the dataset that you want to filter. This can be obtained through the [DescribeDashboardDefinition](#) API operation.

Depending on the column type that you use, there might be restrictions on the types of filters that can be added to an embedded asset. For more information on filter restrictions, see [Filter](#).

- The visual ID of the visual that you want to scope the `FilterGroup` to, if applicable. This can be obtained by using the `getSheetVisuals` method in the Embedding SDK.

In addition to the `getSheetVisuals` method, the `FilterGroup` that you add can only be scoped to the currently selected sheet.

To use this feature, you must already have a dashboard or visual embedded into your application through the Amazon Quick Sight Embedding SDK. For more information about using the Amazon Quick Sight Embedding SDK, see the [amazon-quicksight-embedding-sdk](#) on GitHub.

SDK method interface

Dashboard embedding getter methods

The following table describes different dashboard embedding getter methods that developers can use.

Method	Description
<code>getFilterGroupsForSheet(sheetId: string)</code>	Returns all <code>FilterGroups</code> that are currently scoped to the sheet that is supplied in the parameter.
<code>getFilterGroupsForVisual(sheetId: string, visualId: string)</code>	Returns all <code>FilterGroups</code> that are scoped to the visual that is supplied in the parameter.

If the sheet that is supplied in the parameter is not the currently selected sheet of the embedded dashboard, the above methods return an error.

Visual embedding getter methods

The following table describes different visual embedding getter methods that developers can use.

Method	Description
<code>getFilterGroups()</code>	Returns all <code>FilterGroups</code> that are currently scoped to the embedded visual.

Setter methods

The following table describes different setter methods that developers can use for dashboard or visual embedding.

Method	Description
<code>addFilterGroups(filterGroups: FilterGroup[])</code>	Adds and applies the supplied FilterGroups to the embedded dashboard or visual. A <code>ResponseMessage</code> that indicates whether the addition was successful is returned.
<code>updateFilterGroups(filterGroups: FilterGroup[])</code>	Updates the <code>FilterGroups</code> on the embedded experience that contains the same <code>FilterGroupId</code> as the <code>FilterGroup</code> that is supplied in the parameter. A <code>ResponseMessage</code> that indicates whether the update was successful is returned.
<code>removeFilterGroups(filterGroupsOrIds: FilterGroup[] string[])</code>	Removes the supplied <code>FilterGroups</code> from the dashboard and returns a <code>ResponseMessage</code> that indicates whether the removal attempt is successful.

The `FilterGroup` that is supplied must be scoped to the embedded sheet or visual that is currently selected.

Customize the look and feel of Amazon Quick Sight embedded dashboards and visuals

You can use the Amazon Quick Sight Embedding SDK (version 2.5.0 and higher) to make changes to the theming of your embedded Amazon Quick Sight dashboards and visuals at runtime. Runtime theming makes it easier to integrate your Software as a service (SaaS) application with your Amazon Quick Sight embedded assets. Runtime theming allows you to synchronize the theme of your embedded content with the themes of the parent application that your Amazon Quick Sight assets are embedded into. You can also use runtime theming to add customization options for readers. Theming changes can be applied to embedded assets at initialization or throughout the lifetime of your embedded dashboard or visual.

For more information about themes, see [Using themes in Amazon Quick Sight](#). For more information about using the Amazon Quick Sight Embedding SDK, see the [amazon-quicksight-embedding-sdk](#) on GitHub.

Prerequisites

Before you can get started, make sure that you have the following prerequisites.

- You are using the Amazon Quick Sight Embedding SDK version 2.5.0 or higher.
- Permissions to access the theme that you want to work with. To grant permissions to a theme in Amazon Quick Sight, make a `UpdateThemePermissions` API call or use the **Share** icon next to the theme in the Amazon Quick Sight console's analysis editor.

Terminology and concepts

The following terminology can be useful when working with embedded runtime theming.

- *Theme* – A collection of settings that you can apply to multiple analyses and dashboards that change how the content is displayed.
- *ThemeConfiguration* – A configuration object that contains all of the display properties for a theme.
- *Theme Override* – A *ThemeConfiguration* object that is applied to the active theme to override some or all aspects of how content is displayed.
- *Theme ARN* – An Amazon Resource Name (ARN) that identifies a Amazon Quick Sight theme. Following is an example of custom theme ARN.

```
arn:aws:quicksight:region:account-id:theme/theme-id
```

Amazon Quick Sight provided starter themes don't have a region in their theme ARN. Following is an example of a starter theme ARN.

```
arn:aws:quicksight::aws:theme/CLASSIC
```

Setting up

Make sure that you have the following information ready to get started working with runtime theming.

- The theme ARNs of the themes that you want to use. You can choose an existing theme, or you can create a new one. To obtain a list all themes and theme ARNs in your Amazon Quick Sight account, make a call to the [ListThemes](#) API operation. For information on preset Amazon Quick Sight themes, see [Setting a default theme for Amazon Quick analyses with the Amazon Quick APIs](#).

- If you are using registered user embedding, make sure that the user has access to the themes that you want to use.

If you are using anonymous user embedding, pass a list of theme ARNs to the `AuthorizedResourceArns` parameter of the `GenerateEmbedUrlForAnonymousUser` API. Anonymous users are granted access to any theme that is listed in the `AuthorizedResourceArns` parameter.

SDK method interface

Setter methods

The following table describes different setter methods that developers can use for runtime theming.

Method	Description
<code>setTheme(themeArn: string)</code>	Replaces the active theme of a dashboard or visual with another theme. If applied, the theme override is removed. An error is returned if you don't have access to the theme or if the theme doesn't exist.
<code>setThemeOverride(themeOverride: ThemeConfiguration)</code>	Sets a dynamic <code>ThemeConfiguration</code> to override the current active theme. This replaces the previously set theme override. Any values that are not supplied in the new <code>ThemeConfiguration</code> are defaulted to the values in the currently active theme. An error is returned if the <code>ThemeConfiguration</code> that you supply is invalid.

Initializing embedded content with a theme

To initialize an embedded dashboard or visual with a non-default theme, define a `themeOptions` object on the `DashboardContentOptions` or `VisualContentOptions` parameters, and set the `themeArn` property within `themeOptions` to the desired theme ARN.

The following example initializes an embedded dashboard with the MIDNIGHT theme.

```
import { createEmbeddingContext } from 'amazon-quicksight-embedding-sdk';

const embeddingContext = await createEmbeddingContext();

const {
  embedDashboard,
} = embeddingContext;

const frameOptions = {
  url: '<YOUR_EMBED_URL>',
  container: '#experience-container',
};

const contentOptions = {
  themeOptions: {
    themeArn: "arn:aws:quicksight::aws:theme/MIDNIGHT"
  }
};

// Embedding a dashboard experience
const embeddedDashboardExperience = await embedDashboard(frameOptions, contentOptions);
```

Initializing embedded content with a theme override

Developers can use theme overrides to define the theme of an embedded dashboard or visual at runtime. This allows the dashboard or visual to inherit a theme from a third-party application without the need to pre-configure a theme within Amazon Quick Sight. To initialize an embedded dashboard or visual with a theme override, set the `themeOverride` property within `themeOptions` in either the `DashboardContentOptions` or `VisualContentOptions` parameters. The following example overrides the font of a dashboard's theme from the default font to Amazon Ember.

```
import { createEmbeddingContext } from 'amazon-quicksight-embedding-sdk';

const embeddingContext = await createEmbeddingContext();
```

```
const {
  embedDashboard,
} = embeddingContext;

const frameOptions = {
  url: '<YOUR_EMBED_URL>',
  container: '#experience-container',
};
const contentOptions = {
  themeOptions: {
    "themeOverride":{"Typography":{"FontFamilies":[{"FontFamily":"Comic Neue"]}}}
  }
};

// Embedding a dashboard experience
const embeddedDashboardExperience = await embedDashboard(frameOptions, contentOptions);
```

Initializing embedded content with preloaded themes

Developers can configure a set of dashboard themes to be preloaded on initialization. This is most beneficial for quick toggling between different views, for example dark and light modes. An embedded dashboard or visual can be initialized with up to 5 preloaded themes. To use preloaded themes, set the `preloadThemes` property in either `DashboardContentOptions` or `VisualContentOptions` with an array of up to 5 theme ARNs. The following example preloads the Midnight and Rainier starter themes to a dashboard.

```
import { createEmbeddingContext } from 'amazon-quicksight-embedding-sdk';

const embeddingContext = await createEmbeddingContext();

const {
  embedDashboard,
} = embeddingContext;

const frameOptions = {
  url: '<YOUR_EMBED_URL>',
  container: '#experience-container',
};
const contentOptions = {
  themeOptions: {
    "preloadThemes": ["arn:aws:quicksight::aws:theme/RAINIER",
    "arn:aws:quicksight::aws:theme/MIDNIGHT"]
  }
};
```



```
    }  
};  
  
// Embedding a dashboard experience  
const embeddedDashboardExperience = await embedDashboard(frameOptions, contentOptions);
```

Using the Amazon Quick Sight Embedding SDK to enable shareable links to embedded dashboard views

Amazon Quick Sight developers can use the Amazon Quick Sight Embedding SDK (version 2.8.0 and higher) to allow readers of embedded dashboards to receive and distribute shareable links to their view of an embedded dashboard. Developers can use dashboard or console embedding to generate a shareable link to their application page with Amazon Quick Sight's reference encapsulated using the Amazon Quick Sight Embedding SDK. Amazon Quick Sight Readers can then send this shareable link to their peers. When their peer accesses the shared link, they are taken to the page on the application that contains the embedded Amazon Quick Sight dashboard. Developers can also generate and save shareable links of dashboard views that can be used as a bookmarks for anonymous readers of Amazon Quick Sight when using anonymous embedding.

Prerequisites

Before you get started, make sure that you are using the Amazon Quick Sight Embedding SDK version 2.8.0 or higher

Topics

- [Enabling the SharedView feature configuration for Amazon Quick Sight embedded analytics](#)
- [Creating a shared view with the Amazon Quick Sight createSharedView API](#)
- [Consuming a shared Amazon Quick Sight view](#)

Enabling the SharedView feature configuration for Amazon Quick Sight embedded analytics

When you create an the embedded instance with the Amazon Quick Sight API, set the value of SharedView in the FeatureConfigurations payload to true, as shown in the example below. SharedView overrides the StatePersistence configurations for registered users who access embedded dashboards. If a dashboard user has StatePersistence disabled and SharedView enabled, their state will persist.

```
const generateNewEmbedUrl = async () => {  
    const generateUrlPayload = {
```

```

    experienceConfiguration: {
      QuickSightConsole: {
        FeatureConfigurations: {
          "SharedView": {
            "Enabled": true
          },
        },
      },
    },
  },
}
const result: GenerateEmbedUrlResult = await
generateEmbedUrlForRegisteredUser(generateUrlPayload);
return result.url;
};

```

Creating a shared view with the Amazon Quick Sight createSharedView API

After you update the Embedding SDK to version 2.8.0 or higher, use the `createSharedView` API to create a new shared view. Record the `sharedViewId` and the `dashboardId` that the operation returns. The example below creates a new shared view.

```

const response = await embeddingFrame.createSharedView();
const sharedViewId = response.message.sharedViewId;
const dashboardId = response.message.dashboardId;

```

`createSharedView` can only be called when a user views a dashboard. For console-specific shared view creation, make sure that users are on the dashboard page before you enable the `createSharedView` action. You can do this with the `PAGE_NAVIGATION` event, shown in the example below.

```

const contentOptions = {
  onMessage: async (messageEvent, metadata) => {
    switch (messageEvent.eventName) {
      case 'CONTENT_LOADED': {
        console.log("Do something when the embedded experience is fully
loaded.");
        break;
      }
      case 'ERROR_OCCURRED': {
        console.log("Do something when the embedded experience fails
loading.");
        break;
      }
    }
  }
}

```

```

        case 'PAGE_NAVIGATION': {
            setPageType(messageEvent.message.pageType);
            if (messageEvent.message.pageType === 'DASHBOARD') {
                setShareEnabled(true);
            } else {
                setShareEnabled(false);
            }
            break;
        }
    }
}
};

```

Consuming a shared Amazon Quick Sight view

After you create a new shared view, use the Embedding SDK to make the shared view consumable for other users. The examples below set up a consumable shared view for an embedded dashboard in Amazon Quick Sight.

With an appended URL

Append the `sharedViewId` to the embed URL, under `/views/{viewId}`, and expose this URL to your users. Users can use this URL to will navigate to that shared view.

```

const response = await dashboardFrame.createSharedView();
const newEmbedUrl = await generateNewEmbedUrl();
const formattedUrl = new URL(newEmbedUrl);
formattedUrl.pathname = formattedUrl.pathname.concat('/views/' +
    response.message.sharedViewId);
const baseUrl = formattedUrl.href;
alert("Click to view this QuickSight shared view", baseUrl);

```

With the contentOptions SDK

Pass a `viewId` to the `contentOptions` to open the experience with the given `viewId`.

```

const contentOptions = {
    toolbarOptions: {
        ...
    },
    viewId: sharedViewId,
};

```

```
const embeddedDashboard = await embeddingContext.embedDashboard(
  {container: containerRef.current},
  contentOptions
);
```

With the InitialPath property

```
const shareView = async() => {
  const returnValue = await consoleFrame.createSharedView();
  const {dashboardId, sharedViewId} = returnValue.message;
  const newEmbedUrl = await generateNewEmbedUrl(`/dashboards/${dashboardId}/views/
${sharedViewId}`);
  setShareUrl(newEmbedUrl);
};

const generateNewEmbedUrl = async (initialPath) => {
  const generateUrlPayload = {
    experienceConfiguration: {
      QuickSightConsole: {
        InitialPath: initialPath,
        FeatureConfigurations: {
          "SharedView": {
            "Enabled": true
          },
        },
      },
    },
  }
  const result: GenerateEmbedUrlResult = await
  generateEmbedUrlForRegisteredUser(generateUrlPayload);
  return result.url;
};
```

Embedding Amazon Quick Sight visuals and dashboards with a 1-click embed code

You can embed a visual or dashboard in your application using an embed code. You get this code when you share the dashboard or from the **Embed visual** menu in Amazon Quick Sight.

You can embed a visual or dashboard in your internal application for your registered users. Or you can turn on public sharing in the Amazon Quick Sight console. Doing this grants anyone on the

internet access to a shared visual or dashboard that is embedded in a public application, wiki, or portal.

Following, you can find descriptions about how to embed visuals and dashboards using the 1-click visual or dashboard embed code.

Topics

- [Embedding Amazon Quick Sight visuals and dashboards for registered users with a 1-click embed code](#)
- [Embedding Amazon Quick Sight visuals and dashboards for anonymous users with a 1-click embed code](#)

Embedding Amazon Quick Sight visuals and dashboards for registered users with a 1-click embed code

Applies to: Enterprise Edition

You can embed a visual or dashboard in your internal application for registered users of your Amazon Quick Sight account. You do so using the embed code that you get when you share the dashboard or from the **Embed visual** menu in Amazon Quick Sight. You don't have to run the Amazon Quick Sight embedding API to generate the embed code. You can copy the embed code from Amazon Quick Sight and paste it in your internal application's HTML code.

When users and groups (or all users on your Amazon Quick Sight account) who have access to the dashboard that you want to embed or that holds the visual that you want to embed access your internal application, they're prompted to sign in to the Amazon Quick Sight account with their credentials. After they are authenticated, they can access the visual or dashboard on their internal page. If you have single sign-on enabled, users aren't prompted to sign in again.

Following, you can find descriptions about how to embed a visual or dashboard for registered users using the visual or dashboard embed code.

Before you start

Before you get started, make sure of the following:

- Your internet browser settings contain one of the following to allow communication between the popup and the iframe:

- Native support for the Mozilla Broadcast Channel API. For more information, see [Broadcast Channel API](#) in the Mozilla documentation.
- IndexedDB support.
- LocalStorage support.
- Your internet browser's "block all cookies" settings is turned off.

Step 1: Grant access to the dashboard

For users to access your embedded dashboard, grant them access to view it. You can grant individual users and groups access to a dashboard, or you can grant everyone in your account access. Visual permissions are determined at the dashboard level. To grant access to embedded visuals, grant access to the dashboard that the visual belongs to. For more information, see [Granting access to a dashboard](#).

Step 2: Put the domain where you want to embed the visual or dashboard on your allow list

To embed visuals and dashboards in your internal application, make sure that the domain where you're embedding is allow-listed in your Amazon Quick Sight account. For more information, see [Allow listing static domains](#).

Step 3: Get the embed code

Use the following procedure to get the visual or dashboard embed code.

To get the dashboard embed code

1. Open the published dashboard in Amazon Quick Sight and choose **Share** at upper right. Then choose **Share dashboard**.
2. In the **Share dashboard** page that opens, choose **Copy embed code** at upper left.

The embed code is copied to your clipboard and is similar to the following. The *quicksightdomain* in this example is the URL that you use to access your Amazon Quick Sight account.

```
<iframe
  width="960"
  height="720"
  src="https://quicksightdomain/sn/embed/share/accounts/accountid/
dashboards/dashboardid?directory_alias=account_directory_alias">
```

```
</iframe>
```

To get the visual embed code

1. Open the published dashboard in Amazon Quick Sight and choose the visual that you want to embed. Then open the on-visual menu at the upper right of the visual and choose **Embed visual**.
2. In the **Embed visual** pane that opens, choose **Copy code**.

The embed code is copied to your clipboard and is similar to the following. The *quicksightdomain* in this example is the URL that you use to access your Amazon Quick Sight account.

```
<iframe
  width="600"
  height="400"
  src="https://quicksightdomain/sn/embed/share/accounts/111122223333/
dashboards/DASHBOARDID/sheets/SHEETID>/visuals/VISUALID">
</iframe>
```

Step 4: Paste the code into your internal application's HTML page

Use the following procedure to paste the embed code into your internal application's HTML page

To paste the code in your internal application's HTML page

- Open the HTML code for any page where you want to embed the dashboard and paste the embed code in.

The following example shows what this might look like for an embedded dashboard. The *quicksightdomain* in this example is the URL that you use to access your Amazon Quick Sight account.

```
<!DOCTYPE html>
<html>
<body>

<h2>Example.com - Employee Portal</h2>
<h3>Current shipment stats</h3>
```

```
<iframe
width="960"
height="720"
src="https://quicksightdomain/sn/embed/share/accounts/accountid/
dashboards/dashboardid?directory_alias=account_directory_alias">
</iframe>

</body>
</html>
```

The following example shows what this might look like for an embedded visual. The *quicksightdomain* in this example is the URL that you use to access your Amazon Quick Sight account.

```
<!DOCTYPE html>
<html>
<body>

<h2>Example.com - Employee Portal</h2>
<h3>Current shipment stats</h3>
<iframe
width="600"
height="400"
src="https://quicksightdomain/sn/embed/share/accounts/111122223333/
dashboards/DASHBOARDID/sheets/SHEETID>/visuals/VISUALID?
directory_alias=account_directory_alias">
</iframe>

</body>
</html>
```

For example, let's say that you want to embed your visual or dashboard in an internal Google Sites page. You can open the page on Google Sites and paste the embed code in an embed widget.

If you want to embed your visual or dashboard in an internal Microsoft SharePoint site, you can create a new page and then paste the embed code in an Embed web part.

Embedding Amazon Quick Sight visuals and dashboards for anonymous users with a 1-click embed code

Applies to: Enterprise Edition

You can embed a visual or dashboard in public sites using the embed code that you get when you share the visual or dashboard in Amazon Quick Sight. You can also turn on public sharing by using the Amazon Quick Sight console and automatically grant access to a shared visual or dashboard to anyone on the internet.

Following, you can find how to turn on public sharing for a visual or dashboard and embed the visual or dashboard for anyone on the internet to see. In both cases, you do this by using the 1-click embed code.

Before you start

Before you get started, make sure of the following:

- Your internet browser settings contain one of the following to allow communication between the popup and the iframe that sharing uses:
 - Native support for the Mozilla Broadcast Channel API. For more information, see [Broadcast Channel API](#) in the Mozilla documentation.
 - IndexedDB support.
 - LocalStorage support.
- Your internet browser's "block all cookies" settings is turned off.

Step 1: Turn on public access for the dashboard

For anyone on the internet to access your embedded visual or dashboard, first turn on public access for the dashboard. Visual permissions are determined at the dashboard level. To grant access to embedded visuals, grant access to the dashboard that the visual belongs to. For more information, see [Granting anyone on the internet access to an Amazon Quick Sight dashboard](#).

Step 2: Put the domain where you want to embed the visual or dashboard on your allow list

To embed visuals and dashboards in a public application, wiki, or portal, make sure that the domain where you're embedding it is on the allow list for your Amazon Quick Sight account.

Step 3: Get the embed code

Use the following procedure to get the visual or dashboard embed code.

To get the dashboard embed code

1. Open the published dashboard in Amazon Quick Sight and choose **Share** at upper right. Then choose **Share dashboard**.
2. In the **Share dashboard** page that opens, choose **Copy embed code** at upper left.

The embed code is copied to your clipboard and is similar to the following. The *quicksightdomain* in this example is the URL that you use to access your Amazon Quick Sight account.

```
<iframe
  width="960"
  height="720"
  src="https://quicksightdomain/sn/
    embed/share/accounts/accountid/dashboards/dashboardid">
</iframe>
```

To get the visual embed code

1. Open the published dashboard in Amazon Quick Sight and choose the visual that you want to embed. Then open the on-visual menu in the top right corner of the visual and choose **Embed visual**.
2. In the **Embed visual** pane that opens, choose **Copy code**.

The embed code is copied to your clipboard and is similar to the following. The *quicksightdomain* in this example is the URL that you use to access your Amazon Quick Sight account.

```
<iframe
  width="600"
  height="400"
  src="https://quicksightdomain/sn/embed/share/accounts/111122223333/
dashboards/DASHBOARDID/sheets/SHEETID>/visuals/VISUALID">
</iframe>
```

Step 4: Paste the embed code into an HTML page, wiki page, or portal

Use the following procedure to paste the embed code into an HTML page, wiki page, or portal.

To paste the embed code

- Open the HTML code for the location where you want to embed the visual or dashboard, and paste the embed code in.

The following example shows what this might look like for an embedded dashboard. The *quicksightdomain* in this example is the URL that you use to access your Amazon Quick Sight account.

```
<!DOCTYPE html>
<html>
<body>

<h2>Example.com - Employee Portal</h2>
<h3>Current shipment stats</h3>
  <iframe
    width="960"
    height="720"
    src="https://quicksightdomain/sn/
      embed/share/accounts/accountid/dashboards/dashboardid">
  </iframe>

</body>
</html>
```

The following example shows what this might look like for an embedded visual. The *quicksightdomain* in this example is the URL that you use to access your Amazon Quick Sight account.

```
<!DOCTYPE html>
<html>
<body>

<h2>Example.com - Employee Portal</h2>
<h3>Current shipment stats</h3>
  <iframe
    width="600"
    height="400"
```

```
src="https://quicksightdomain/sn/embed/share/accounts/111122223333/
dashboards/DASHBOARDID/sheets/SHEETID>/visuals/VISUALID">
</iframe>

</body>
</html>
```

If your public-facing applications are built on Google Sites, open the page on Google Sites and then paste the embed code using the embed widget.

Make sure that the following domains in Amazon Quick Sight are on your allow list when you embed in Google Sites:

- <https://googleusercontent.com> (turns on subdomains)
- <https://www.gstatic.com>
- <https://sites.google.com>

After you embed the visual or dashboard in your application, anyone who can access your application can access the embedded visual or dashboard. To update a dashboard that's shared with the public, see [Updating a publicly shared dashboard](#). To turn off public sharing, see [Turning off public sharing settings](#).

When you turn off public sharing, no one from the internet can access a dashboard or dashboards that you have embedded on a public application or shared with a link. The next time anyone tries to view such a dashboard from the internet, they receive a message that they don't have access to view it.

Embedding with the Amazon Quick Sight APIs

Applies to: Enterprise Edition

Intended audience: Amazon Quick developers

There are only a few steps involved in the actual process of embedding analytics using the Amazon Quick Sight APIs.

Before you begin, make sure to have the following items in place:

- Set up the required IAM permissions for the caller identity used by your application that will use the AWS SDK to make API calls. For example, grant permission to allow the `quicksight:GenerateEmbedUrlForAnonymousUser` or `quicksight:GenerateEmbedUrlForRegisteredUser` action.
- To embed for registered users, share Amazon Quick Sight assets with them beforehand. For new authenticating users, know how to grant access to the assets. One way to do this is by adding all the assets to a Amazon Quick Sight folder. If you prefer to use the Amazon Quick Sight API, use the `DescribeDashboardPermissions` and `UpdateDashboardPermissions` API operations. For more information, see [DescribeDashboardPermissions](#) or [UpdateDashboardPermissions](#) in the *Amazon Quick API Reference*. If you want to share the dashboard with all users in a namespace or group, you can share the dashboard with namespace or group.
- If you're embedding dashboards, make sure to have the ID of the dashboards you want to embed. The dashboard ID is the code in the URL of the dashboard. You can also get it from the dashboard URL.
- A Amazon Quick Sight administrator must explicitly enable domains where you plan to embed your Amazon Quick Sight analytics. You can do this by using the **Manage Amazon Quick Sight, Domains and Embedding** from the profile menu, or you can use the `AllowedDomains` parameter of a `GenerateEmbedUrlForAnonymousUser` or `GenerateEmbedUrlForRegisteredUser` API call.

This option is only visible to Amazon Quick Sight administrators. You can also add subdomains as part of a domain. For more information, see [Allow listing domains at runtime with the Amazon Quick API](#).

All domains in your static allow list (such as development, staging, and production) must be explicitly allowed, and they must use HTTPS. You can add up to 100 domains to the allow list. You can add domains at runtime with Amazon Quick Sight API operations.

After all the prerequisites are complete, embedding Amazon Quick Sight involves the following steps, which are explained in greater detail later:

1. For authentication, use your application server to authenticate the user. After authentication in your server, generate the embedded dashboard URL using the AWS SDK that you need.
2. In your web portal or application, embed Amazon Quick Sight using the generated URL. To simplify this process, you can use the Amazon Quick Sight Embedding SDK, available on [NPMJS](#)

and [GitHub](#). This customized JavaScript SDK is designed to help you efficiently integrate Amazon Quick Sight into your application pages, set defaults, connect controls, get callbacks, and handle errors.

You can use AWS CloudTrail auditing logs to get information about the number of embedded dashboards, users of an embedded experience, and access rates.

Topics

- [Embedding Amazon Quick Sight dashboards with the Amazon Quick Sight API](#)
- [Embedding Amazon Quick Sight visuals with the Amazon Quick Sight APIs](#)
- [Embedding the full functionality of the Amazon Quick Sight console for registered users](#)
- [Embedding the Amazon Q in Amazon Quick Sight Generative Q&A experience](#)
- [Embedding the Amazon Quick Sight Q search bar \(Classic\)](#)
- [Embedding analytics using the GetDashboardEmbedURL and GetSessionEmbedURL API operations](#)

Embedding Amazon Quick Sight dashboards with the Amazon Quick Sight API

Use the following topics to learn about embedding dashboards with the Amazon Quick Sight API.

Topics

- [Embedding Amazon Quick Sight dashboards for registered users](#)
- [Embedding Amazon Quick Sight dashboards for anonymous \(unregistered\) users](#)
- [Enabling executive summaries in embedded dashboards](#)

Embedding Amazon Quick Sight dashboards for registered users

Important

Amazon Quick Sight has new API operations for embedding analytics: `GenerateEmbedUrlForAnonymousUser` and `GenerateEmbedUrlForRegisteredUser`. You can still use the `GetDashboardEmbedUrl` and `GetSessionEmbedUrl` API operations to embed dashboards and the Amazon Quick Sight console, but they don't contain the latest embedding capabilities. For more information about embedding using the

old API operations, see [Embedding analytics using the GetDashboardEmbedURL and GetSessionEmbedURL API operations](#).

Applies to: Enterprise Edition

Intended audience: Amazon Quick developers

In the following sections, you can find detailed information about how to set up embedded Amazon Quick Sight dashboards for registered users of Amazon Quick Sight.

Topics

- [Step 1: Set up permissions](#)
- [Step 2: Generate the URL with the authentication code attached](#)
- [Step 3: Embed the dashboard URL](#)

Step 1: Set up permissions

In the following section, you can find out how to set up permissions for the backend application or web server. This task requires administrative access to IAM.

Each user who accesses a dashboard assumes a role that gives them Amazon Quick Sight access and permissions to the dashboard. To make this possible, create an IAM role in your AWS account. Associate an IAM policy with the role to provide permissions to any user who assumes it. The IAM role needs to provide permissions to retrieve embedding URLs for a specific user pool. With the help of the wildcard character *, you can grant the permissions to generate a URL for all users in a specific namespace, or for a subset of users in specific namespaces. For this, you add `quicksight:GenerateEmbedUrlForRegisteredUser`.

You can create a condition in your IAM policy that limits the domains that developers can list in the `AllowedDomains` parameter of a `GenerateEmbedUrlForRegisteredUser` API operation. The `AllowedDomains` parameter is an optional parameter. It grants you as a developer the option to override the static domains that are configured in the **Manage Amazon Quick Sight** menu. Instead, you can list up to three domains or subdomains that can access the generated URL. This URL is then embedded in the website that you create. Only the domains that are listed in the parameter

can access the embedded visual. Without this condition, you can list any domain on the internet in the `AllowedDomains` parameter.

To limit the domains that developers can use with this parameter, add an `AllowedEmbeddingDomains` condition to your IAM policy. For more information about the `AllowedDomains` parameter, see [GenerateEmbedUrlForRegisteredUser](#) in the *Amazon Quick Sight API Reference*.

Security best practice for IAM condition operators

Improperly configured IAM condition operators can allow unauthorized access to your embedded Quick resources through URL variations. When using the `quicksight:AllowedEmbeddingDomains` condition key in your IAM policies, use condition operators that either allow specific domains or deny all domains that are not specifically allowed. For more information about IAM condition operators, see [IAM JSON policy elements: Condition operators](#) in the IAM User Guide.

Many different URL variations can point to the same resource. For example, the following URLs all resolve to the same content:

- `https://example.com`
- `https://example.com/`
- `https://Example.com`

If your policy uses operators that do not account for these URL variations, an attacker can bypass your restrictions by providing equivalent URL variations.

You must validate that your IAM policy uses appropriate condition operators to prevent bypass vulnerabilities and ensure that only your intended domains can access your embedded resources.

The following sample policy provides these permissions.

Additionally, if you are creating first-time users who will be Amazon Quick Sight readers, make sure to add the `quicksight:RegisterUser` permission in the policy.

The following sample policy provides permission to retrieve an embedding URL for first-time users who are to be Amazon Quick Sight readers.

Finally, your application's IAM identity must have a trust policy associated with it to allow access to the role that you just created. This means that when a user accesses your application, your application can assume the role on the user's behalf and provision the user in Amazon Quick Sight. The following example shows a sample trust policy.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowLambdaFunctionsToAssumeThisRole",
      "Effect": "Allow",
      "Principal": {
        "Service": "lambda.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    },
    {
      "Sid": "AllowEC2InstancesToAssumeThisRole",
      "Effect": "Allow",
      "Principal": {
        "Service": "ec2.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

For more information regarding trust policies for OpenID Connect or SAML authentication, see the following sections of the *IAM User Guide*:

- [Creating a Role for Web Identity or OpenID Connect Federation \(Console\)](#)
- [Creating a Role for SAML 2.0 Federation \(Console\)](#)

Step 2: Generate the URL with the authentication code attached

In the following section, you can find out how to authenticate your user and get the embeddable dashboard URL on your application server. If you plan to embed dashboards for IAM or Amazon Quick Sight identity types, share the dashboard with the users.

When a user accesses your app, the app assumes the IAM role on the user's behalf. Then it adds the user to Amazon Quick Sight, if that user doesn't already exist. Next, it passes an identifier as the unique role session ID.

Performing these steps ensures that each viewer of the dashboard is uniquely provisioned in Amazon Quick Sight. It also enforces per-user settings, such as the row-level security and dynamic defaults for parameters.

The following examples perform the IAM authentication on the user's behalf. This code runs on your app server.

Java

```
import com.amazonaws.auth.AWSCredentials;
import com.amazonaws.auth.BasicAWSCredentials;
import com.amazonaws.auth.AWSCredentialsProvider;
import com.amazonaws.regions.Regions;
import com.amazonaws.services.quicksight.AmazonQuickSight;
import com.amazonaws.services.quicksight.AmazonQuickSightClientBuilder;
import
com.amazonaws.services.quicksight.model.GenerateEmbedUrlForRegisteredUserRequest;
import
com.amazonaws.services.quicksight.model.GenerateEmbedUrlForRegisteredUserResult;
import
com.amazonaws.services.quicksight.model.RegisteredUserEmbeddingExperienceConfiguration;
import
com.amazonaws.services.quicksight.model.RegisteredUserDashboardEmbeddingConfiguration;

/**
 * Class to call QuickSight AWS SDK to get url for dashboard embedding.
 */
public class GetQuicksightEmbedUrlRegisteredUserDashboardEmbedding {

    private final AmazonQuickSight quickSightClient;

    public GetQuicksightEmbedUrlRegisteredUserDashboardEmbedding() {
        this.quickSightClient = AmazonQuickSightClientBuilder
```

```

        .standard()
        .withRegion(Regions.US_EAST_1.getName())
        .withCredentials(new AWSCredentialsProvider() {
            @Override
            public AWSCredentials getCredentials() {
                // provide actual IAM access key and secret key here
                return new BasicAWSCredentials("access-key", "secret-key");
            }

            @Override
            public void refresh() {}
        })
        .build();
    }

    public String getQuicksightEmbedUrl(
        final String accountId, // AWS Account ID
        final String dashboardId, // Dashboard ID to embed
        final List<String> allowedDomains, // Runtime allowed domain for
embedding
        final String userArn // Registered user arn to use for embedding. Refer
to Get Embed Url section in developer portal to find out how to get user arn for a
QuickSight user.
    ) throws Exception {
        final RegisteredUserEmbeddingExperienceConfiguration
experienceConfiguration = new RegisteredUserEmbeddingExperienceConfiguration()
            .withDashboard(new
RegisteredUserDashboardEmbeddingConfiguration().withInitialDashboardId(dashboardId));
        final GenerateEmbedUrlForRegisteredUserRequest
generateEmbedUrlForRegisteredUserRequest = new
GenerateEmbedUrlForRegisteredUserRequest();
        generateEmbedUrlForRegisteredUserRequest.setAwsAccountId(accountId);
        generateEmbedUrlForRegisteredUserRequest.setUserArn(userArn);
        generateEmbedUrlForRegisteredUserRequest.setAllowedDomains(allowedDomains);

generateEmbedUrlForRegisteredUserRequest.setExperienceConfiguration(experienceConfiguration);

        final GenerateEmbedUrlForRegisteredUserResult
generateEmbedUrlForRegisteredUserResult =
quicksightClient.generateEmbedUrlForRegisteredUser(generateEmbedUrlForRegisteredUserRequest);

        return generateEmbedUrlForRegisteredUserResult.getEmbedUrl();
    }

```

```
}
```

JavaScript

```
global.fetch = require('node-fetch');
const AWS = require('aws-sdk');

function generateEmbedUrlForRegisteredUser(
    accountId,
    dashboardId,
    openIdToken, // Cognito-based token
    userArn, // registered user arn
    roleArn, // IAM user role to use for embedding
    sessionName, // Session name for the roleArn assume role
    allowedDomains, // Runtime allowed domain for embedding
    getEmbedUrlCallback, // GetEmbedUrl success callback method
    errorCallback // GetEmbedUrl error callback method
) {
    const stsClient = new AWS.STS();
    let stsParams = {
        RoleSessionName: sessionName,
        WebIdentityToken: openIdToken,
        RoleArn: roleArn
    }

    stsClient.assumeRoleWithWebIdentity(stsParams, function(err, data) {
        if (err) {
            console.log('Error assuming role');
            console.log(err, err.stack);
            errorCallback(err);
        } else {
            const getDashboardParams = {
                "AwsAccountId": accountId,
                "ExperienceConfiguration": {
                    "Dashboard": {
                        "InitialDashboardId": dashboardId
                    }
                },
                "UserArn": userArn,
                "AllowedDomains": allowedDomains,
                "SessionLifetimeInMinutes": 600
            };
        }
    });
}
```

```

        const quicksightClient = new AWS.QuickSight({
            region: process.env.AWS_REGION,
            credentials: {
                accessKeyId: data.Credentials.AccessKeyId,
                secretAccessKey: data.Credentials.SecretAccessKey,
                sessionToken: data.Credentials.SessionToken,
                expiration: data.Credentials.Expiration
            }
        });

        quicksightClient.generateEmbedUrlForRegisteredUser(getDashboardParams,
function(err, data) {
    if (err) {
        console.log(err, err.stack);
        errorCallback(err);
    } else {
        const result = {
            "statusCode": 200,
            "headers": {
                "Access-Control-Allow-Origin": "*", // Use your website
domain to secure access to GetEmbedUrl API
                "Access-Control-Allow-Headers": "Content-Type"
            },
            "body": JSON.stringify(data),
            "isBase64Encoded": false
        }
        getEmbedUrlCallback(result);
    }
});
    }
});
}

```

Python3

```

import json
import boto3
from botocore.exceptions import ClientError

sts = boto3.client('sts')

# Function to generate embedded URL
# accountId: AWS account ID

```

```

# dashboardId: Dashboard ID to embed
# userArn: arn of registered user
# allowedDomains: Runtime allowed domain for embedding
# roleArn: IAM user role to use for embedding
# sessionName: session name for the roleArn assume role
def getEmbeddingURL(accountId, dashboardId, userArn, allowedDomains, roleArn,
    sessionName):
    try:
        assumedRole = sts.assume_role(
            RoleArn = roleArn,
            RoleSessionName = sessionName,
        )
    except ClientError as e:
        return "Error assuming role: " + str(e)
    else:
        assumedRoleSession = boto3.Session(
            aws_access_key_id = assumedRole['Credentials']['AccessKeyId'],
            aws_secret_access_key = assumedRole['Credentials']['SecretAccessKey'],
            aws_session_token = assumedRole['Credentials']['SessionToken'],
        )
        try:
            quicksightClient = assumedRoleSession.client('quicksight', region_name='us-
west-2')
            response = quicksightClient.generate_embed_url_for_registered_user(
                AwsAccountId=accountId,
                ExperienceConfiguration = {
                    "Dashboard": {
                        "InitialDashboardId": dashboardId
                    }
                },
                UserArn = userArn,
                AllowedDomains = allowedDomains,
                SessionLifetimeInMinutes = 600
            )

            return {
                'statusCode': 200,
                'headers': {"Access-Control-Allow-Origin": "*", "Access-Control-Allow-
Headers": "Content-Type"},
                'body': json.dumps(response),
                'isBase64Encoded': bool('false')
            }
        except ClientError as e:

```

```
return "Error generating embedding url: " + str(e)
```

Node.js

The following example shows the JavaScript (Node.js) that you can use on the app server to generate the URL for the embedded dashboard. You can use this URL in your website or app to display the dashboard.

Example

```
const AWS = require('aws-sdk');
const https = require('https');

var quicksightClient = new AWS.Service({
  apiConfig: require('./quicksight-2018-04-01.min.json'),
  region: 'us-east-1',
});

quicksightClient.generateEmbedUrlForRegisteredUser({
  'AwsAccountId': '111122223333',
  'ExperienceConfiguration': {
    'Dashboard': {
      'InitialDashboardId': '1c1fe111-e2d2-3b30-44ef-a0e111111cde'
    }
  },
  'UserArn': 'REGISTERED_USER_ARN',
  'AllowedDomains': allowedDomains,
  'SessionLifetimeInMinutes': 100
}, function(err, data) {
  console.log('Errors: ');
  console.log(err);
  console.log('Response: ');
  console.log(data);
});
```

Example

```
//The URL returned is over 900 characters. For this example, we've shortened the string
for
  //readability and added ellipsis to indicate that it's incomplete.
  {
    Status: 200,
```

```
        EmbedUrl: 'https://quicksightdomain/embed/12345/dashboards/67890...'
        RequestId: '7bee030e-f191-45c4-97fe-d9faf0e03713'
    }
```

.NET/C#

The following example shows the .NET/C# code that you can use on the app server to generate the URL for the embedded dashboard. You can use this URL in your website or app to display the dashboard.

Example

```
using System;
using Amazon.QuickSight;
using Amazon.QuickSight.Model;

namespace GenerateDashboardEmbedUrlForRegisteredUser
{
    class Program
    {
        static void Main(string[] args)
        {
            var quicksightClient = new AmazonQuickSightClient(
                AccessKey,
                SecretAccessKey,
                SessionToken,
                Amazon.RegionEndpoint.USEast1);
            try
            {
                RegisteredUserDashboardEmbeddingConfiguration
registeredUserDashboardEmbeddingConfiguration
                = new RegisteredUserDashboardEmbeddingConfiguration
                {
                    InitialDashboardId = "dashboardId"
                };
                RegisteredUserEmbeddingExperienceConfiguration
registeredUserEmbeddingExperienceConfiguration
                = new RegisteredUserEmbeddingExperienceConfiguration
                {
                    Dashboard = registeredUserDashboardEmbeddingConfiguration
                };

                Console.WriteLine(
```



```

        quicksightClient.GenerateEmbedUrlForRegisteredUserAsync(new
GenerateEmbedUrlForRegisteredUserRequest
        {
            AwsAccountId = "111122223333",
            ExperienceConfiguration =
registeredUserEmbeddingExperienceConfiguration,
            UserArn = "REGISTERED_USER_ARN",
            AllowedDomains = allowedDomains,
            SessionLifetimeInMinutes = 100
        }).Result.EmbedUrl
    );
} catch (Exception ex) {
    Console.WriteLine(ex.Message);
}
}
}
}

```

AWS CLI

To assume the role, choose one of the following AWS Security Token Service (AWS STS) API operations:

- [AssumeRole](#) – Use this operation when you're using an IAM identity to assume the role.
- [AssumeRoleWithWebIdentity](#) – Use this operation when you're using a web identity provider to authenticate your user.
- [AssumeRoleWithSaml](#) – Use this operation when you're using SAML to authenticate your users.

The following example shows the CLI command to set the IAM role. The role needs to have permissions enabled for `quicksight:GenerateEmbedUrlForRegisteredUser`. If you are taking a just-in-time approach to add users when they first open a dashboard, the role also needs permissions enabled for `quicksight:RegisterUser`.

```

aws sts assume-role \
    --role-arn "arn:aws:iam::111122223333:role/embedding_quicksight_dashboard_role"
\
    --role-session-name john.doe@example.com

```

The `assume-role` operation returns three output parameters: the access key, the secret key, and the session token.

Note

If you get an `ExpiredToken` error when calling the `AssumeRole` operation, this is probably because the previous `SESSION_TOKEN` is still in the environment variables. Clear this by setting the following variables:

- `AWS_ACCESS_KEY_ID`
- `AWS_SECRET_ACCESS_KEY`
- `AWS_SESSION_TOKEN`

The following example shows how to set these three parameters in the CLI. If you're using a Microsoft Windows machine, use `set` instead of `export`.

```
export AWS_ACCESS_KEY_ID      = "access_key_from_assume_role"
export AWS_SECRET_ACCESS_KEY  = "secret_key_from_assume_role"
export AWS_SESSION_TOKEN      = "session_token_from_assume_role"
```

Running these commands sets the role session ID of the user visiting your website to `embedding_quicksight_dashboard_role/john.doe@example.com`. The role session ID is made up of the role name from `role-arn` and the `role-session-name` value. Using the unique role session ID for each user ensures that appropriate permissions are set for each user. It also prevents any throttling of user access. *Throttling* is a security feature that prevents the same user from accessing Amazon Quick Sight from multiple locations.

The role session ID also becomes the user name in Amazon Quick Sight. You can use this pattern to provision your users in Amazon Quick Sight ahead of time, or to provision them the first time they access the dashboard.

The following example shows the CLI command that you can use to provision a user. For more information about [RegisterUser](#), [DescribeUser](#), and other Amazon Quick Sight API operations, see the [Amazon Quick Sight API Reference](#).

```
aws quicksight register-user \
  --aws-account-id 111122223333 \
  --namespace default \
  --identity-type IAM \
  --iam-arn "arn:aws:iam::111122223333:role/embedding_quicksight_dashboard_role" \
```

```
--user-role READER \
--user-name jhnd \
--session-name "john.doe@example.com" \
--email john.doe@example.com \
--region us-east-1 \
--custom-permissions-name TeamA1
```

If your user is authenticated through Microsoft AD, you don't need to use `RegisterUser` to set them up. Instead, they should be automatically subscribed the first time they access Amazon Quick Sight. For Microsoft AD users, you can use `DescribeUser` to get the user ARN.

The first time a user accesses Amazon Quick Sight, you can also add this user to the group that the dashboard is shared with. The following example shows the CLI command to add a user to a group.

```
aws quicksight create-group-membership \
  --aws-account-id=111122223333 \
  --namespace=default \
  --group-name=financeusers \
  --member-name="embedding_quicksight_dashboard_role/john.doe@example.com"
```

You now have a user of your app who is also a user of Amazon Quick Sight, and who has access to the dashboard.

Finally, to get a signed URL for the dashboard, call `generate-embed-url-for-registered-user` from the app server. This returns the embeddable dashboard URL. The following example shows how to generate the URL for an embedded dashboard using a server-side call for users authenticated through AWS Managed Microsoft AD or single sign-on (IAM Identity Center).

```
aws quicksight generate-embed-url-for-registered-user \
  --aws-account-id 111122223333 \
  --session-lifetime-in-minutes 600 \
  --user-arn arn:aws:quicksight:us-east-1:111122223333:user/default/  
embedding_quicksight_visual_role/embeddingsession \
  --allowed-domains '["domain1", "domain2"]' \
  --experience-configuration  
Dashboard={InitialDashboardId=1a1ac2b2-3fc3-4b44-5e5d-c6db6778df89}
```

For more information about using this operation, see [GenerateEmbedUrlForRegisteredUser](#). You can use this and other API operations in your own code.

Step 3: Embed the dashboard URL

In the following section, you can find out how you can use the [Amazon Quick Sight Embedding SDK](#) (JavaScript) to embed the dashboard URL from step 3 in your website or application page. With the SDK, you can do the following:

- Place the dashboard on an HTML page.
- Pass parameters into the dashboard.
- Handle error states with messages that are customized to your application.

Call the `GenerateEmbedUrlForRegisteredUser` API operation to generate the URL that you can embed in your app. This URL is valid for 5 minutes, and the resulting session is valid for up to 10 hours. The API operation provides the URL with an `auth_code` that enables a single-sign on session.

The following shows an example response from `generate-embed-url-for-registered-user`.

```
//The URL returned is over 900 characters. For this example, we've shortened the string
for
//readability and added ellipsis to indicate that it's incomplete.
{
  "Status": "200",
  "EmbedUrl": "https://quicksightdomain/embed/12345/dashboards/67890..",
  "RequestId": "7bee030e-f191-45c4-97fe-d9faf0e03713"
}
```

Embed this dashboard in your webpage by using the [Amazon Quick Sight Embedding SDK](#) or by adding this URL into an `iframe`. If you set a fixed height and width number (in pixels), Amazon Quick Sight uses those and doesn't change your visual as your window resizes. If you set a relative percent height and width, Amazon Quick Sight provides a responsive layout that is modified as your window size changes. By using the Amazon Quick Sight Embedding SDK, you can also control parameters within the dashboard and receive callbacks in terms of page load completion and errors.

The domain that is going to host embedded dashboards must be on the *allow list*, the list of approved domains for your Quick subscription. This requirement protects your data by keeping unapproved domains from hosting embedded dashboards. For more information about adding domains for embedded dashboards, see [Allow listing domains at runtime with the Amazon Quick Sight API](#).

The following example shows how to use the generated URL. This code is generated on your app server.

SDK 2.0

```
<!DOCTYPE html>
<html>

  <head>
    <title>Dashboard Embedding Example</title>
    <script src="https://unpkg.com/amazon-quicksight-embedding-sdk@2.0.0/dist/
quicksight-embedding-js-sdk.min.js"></script>
    <script type="text/javascript">
      const embedDashboard = async() => {
        const {
          createEmbeddingContext,
        } = QuickSightEmbedding;

        const embeddingContext = await createEmbeddingContext({
          onChange: (changeEvent, metadata) => {
            console.log('Context received a change', changeEvent,
metadata);
          },
        });

        const frameOptions = {
          url: '<YOUR_EMBED_URL>',
          container: '#experience-container',
          height: "700px",
          width: "1000px",
          onChange: (changeEvent, metadata) => {
            switch (changeEvent.eventName) {
              case 'FRAME_MOUNTED': {
                console.log("Do something when the experience frame is
mounted.");
                break;
              }
              case 'FRAME_LOADED': {
                console.log("Do something when the experience frame is
loaded.");
                break;
              }
            }
          },
        },
```

```

};

const contentOptions = {
  parameters: [
    {
      Name: 'country',
      Values: [
        'United States'
      ],
    },
    {
      Name: 'states',
      Values: [
        'California',
        'Washington'
      ]
    }
  ],
  locale: "en-US",
  sheetOptions: {
    initialSheetId: '<YOUR_SHEETID>',
    singleSheet: false,
    emitSizeChangedEventOnSheetChange: false,
  },
  toolbarOptions: {
    export: false,
    undoRedo: false,
    reset: false
  },
  attributionOptions: {
    overlayContent: false,
  },
  onMessage: async (messageEvent, experienceMetadata) => {
    switch (messageEvent.eventName) {
      case 'CONTENT_LOADED': {
        console.log("All visuals are loaded. The title of the
document:", messageEvent.message.title);
        break;
      }
      case 'ERROR_OCCURRED': {
        console.log("Error occurred while rendering the
experience. Error code:", messageEvent.message.errorCode);
        break;
      }
    }
  }
}

```

```

        case 'PARAMETERS_CHANGED': {
            console.log("Parameters changed. Changed parameters:",
messageEvent.message.changedParameters);
            break;
        }
        case 'SELECTED_SHEET_CHANGED': {
            console.log("Selected sheet changed. Selected sheet:",
messageEvent.message.selectedSheet);
            break;
        }
        case 'SIZE_CHANGED': {
            console.log("Size changed. New dimensions:",
messageEvent.message);
            break;
        }
        case 'MODAL_OPENED': {
            window.scrollTo({
                top: 0 // iframe top position
            });
            break;
        }
    }
},
    ];
    const embeddedDashboardExperience = await
embeddingContext.embedDashboard(frameOptions, contentOptions);

    const selectCountryElement = document.getElementById('country');
    selectCountryElement.addEventListener('change', (event) => {
        embeddedDashboardExperience.setParameters([
            {
                Name: 'country',
                Values: event.target.value
            }
        ]);
    });
};
</script>
</head>

<body onload="embedDashboard()">
    <span>
        <label for="country">Country</label>
        <select id="country" name="country">

```

```

        <option value="United States">United States</option>
        <option value="Mexico">Mexico</option>
        <option value="Canada">Canada</option>
    </select>
</span>
<div id="experience-container"></div>
</body>

</html>

```

SDK 1.0

```

<!DOCTYPE html>
<html>

    <head>
        <title>Basic Embed</title>
        <script src="https://unpkg.com/amazon-quicksight-embedding-sdk@1.0.15/dist/
quicksight-embedding-js-sdk.min.js"></script>
        <script type="text/javascript">
            var dashboard
            function onDashboardLoad(payload) {
                console.log("Do something when the dashboard is fully loaded.");
            }

            function onError(payload) {
                console.log("Do something when the dashboard fails loading");
            }

            function embedDashboard() {
                var containerDiv = document.getElementById("embeddingContainer");
                var options = {
                    // replace this dummy url with the one generated via embedding API
                    url: "https://us-east-1.quicksight.aws.amazon.com/sn/dashboards/
dashboardId?isauthcode=true&identityprovider=quicksight&code=authcode",
                    container: containerDiv,
                    parameters: {
                        country: "United States"
                    },
                    scrolling: "no",
                    height: "700px",
                    width: "1000px",
                    locale: "en-US",

```



```
        footerPaddingEnabled: true
    };
    dashboard = QuickSightEmbedding.embedDashboard(options);
    dashboard.on("error", onError);
    dashboard.on("load", onDashboardLoad);
}

function onCountryChange(obj) {
    dashboard.setParameters({country: obj.value});
}
</script>
</head>

<body onload="embedDashboard()">
    <span>
        <label for="country">Country</label>
        <select id="country" name="country" onchange="onCountryChange(this)">
            <option value="United States">United States</option>
            <option value="Mexico">Mexico</option>
            <option value="Canada">Canada</option>
        </select>
    </span>
    <div id="embeddingContainer"></div>
</body>

</html>
```

For this example to work, make sure to use the Amazon Quick Sight Embedding SDK to load the embedded dashboard on your website using JavaScript. To get your copy, do one of the following:

- Download the [Amazon Quick Sight Embedding SDK](#) from GitHub. This repository is maintained by a group of Amazon Quick Sight developers.
- Download the latest embedding SDK version from <https://www.npmjs.com/package/amazon-quicksight-embedding-sdk>.
- If you use npm for JavaScript dependencies, download and install it by running the following command.

```
npm install amazon-quicksight-embedding-sdk
```

Embedding Amazon Quick Sight dashboards for anonymous (unregistered) users

Important

Amazon Quick Sight has new API operations for embedding analytics: `GenerateEmbedUrlForAnonymousUser` and `GenerateEmbedUrlForRegisteredUser`. You can still use the `GetDashboardEmbedUrl` and `GetSessionEmbedUrl` API operations to embed dashboards and the Amazon Quick Sight console, but they don't contain the latest embedding capabilities. For more information about embedding using the old API operations, see [Embedding analytics using the `GetDashboardEmbedURL` and `GetSessionEmbedURL` API operations](#).

Applies to: Enterprise Edition

Intended audience: Amazon Quick developers

In the following sections, you can find detailed information about how to set up embedded Amazon Quick Sight dashboards for anonymous (unregistered) users.

Topics

- [Step 1: Set up permissions](#)
- [Step 2: Generate the URL with the authentication code attached](#)
- [Step 3: Embed the dashboard URL](#)

Step 1: Set up permissions

Applies to: Enterprise Edition

Intended audience: Amazon Quick developers

In the following section, you can find out how to set up permissions for the backend application or web server. This task requires administrative access to IAM.

Each user who accesses a dashboard assumes a role that gives them Amazon Quick Sight access and permissions to the dashboard. To make this possible, create an IAM role in your AWS account. Associate an IAM policy with the role to provide permissions to any user who assumes it.

You can create a condition in your IAM policy that limits the domains that developers can list in the `AllowedDomains` parameter of a `GenerateEmbedUrlForAnonymousUser` API operation. The `AllowedDomains` parameter is an optional parameter. It grants you as a developer the option to override the static domains that are configured in the **Manage Amazon Quick Sight** menu. Instead, you can list up to three domains or subdomains that can access a generated URL. This URL is then embedded in the website that you create. Only the domains that are listed in the parameter can access the embedded dashboard. Without this condition, you can list any domain on the internet in the `AllowedDomains` parameter.

To limit the domains that developers can use with this parameter, add an `AllowedEmbeddingDomains` condition to your IAM policy. For more information about the `AllowedDomains` parameter, see [GenerateEmbedUrlForAnonymousUser](#) in the *Amazon Quick Sight API Reference*.

Security best practice for IAM condition operators

Improperly configured IAM condition operators can allow unauthorized access to your embedded Quick resources through URL variations. When using the `quicksight:AllowedEmbeddingDomains` condition key in your IAM policies, use condition operators that either allow specific domains or deny all domains that are not specifically allowed. For more information about IAM condition operators, see [IAM JSON policy elements: Condition operators](#) in the IAM User Guide.

Many different URL variations can point to the same resource. For example, the following URLs all resolve to the same content:

- `https://example.com`
- `https://example.com/`
- `https://Example.com`

If your policy uses operators that do not account for these URL variations, an attacker can bypass your restrictions by providing equivalent URL variations.

You must validate that your IAM policy uses appropriate condition operators to prevent bypass vulnerabilities and ensure that only your intended domains can access your embedded resources.

The following sample policy provides these permissions for use with `GenerateEmbedUrlForAnonymousUser`. For this approach to work, you also need a session pack, or session capacity pricing, for your AWS account. Otherwise, when a user tries to access the dashboard, the error `UnsupportedPricingPlanException` is returned.

Your application's IAM identity must have a trust policy associated with it to allow access to the role that you just created. This means that when a user accesses your application, your application can assume the role on the user's behalf to open the dashboard. The following example shows a sample trust policy.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowLambdaFunctionsToAssumeThisRole",
      "Effect": "Allow",
      "Principal": {
        "Service": "lambda.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    },
    {
      "Sid": "AllowEC2InstancesToAssumeThisRole",
      "Effect": "Allow",
      "Principal": {
        "Service": "ec2.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

For more information regarding trust policies, see [Temporary security credentials in IAM](#) in the *IAM User Guide*.

Step 2: Generate the URL with the authentication code attached

Applies to: Enterprise Edition

Intended audience: Amazon Quick developers

In the following section, you can find how to authenticate on behalf of the anonymous visitor and get the embeddable dashboard URL on your application server.

When a user accesses your app, the app assumes the IAM role on the user's behalf. Then it adds the user to Amazon Quick Sight, if that user doesn't already exist. Next, it passes an identifier as the unique role session ID.

The following examples perform the IAM authentication on the user's behalf. It passes an identifier as the unique role session ID. This code runs on your app server.

Java

```
import java.util.List;
import com.amazonaws.auth.AWSCredentials;
import com.amazonaws.auth.AWSCredentialsProvider;
import com.amazonaws.auth.BasicAWSCredentials;
import com.amazonaws.regions.Regions;
import com.amazonaws.services.quicksight.AmazonQuickSight;
import com.amazonaws.services.quicksight.AmazonQuickSightClientBuilder;
import
com.amazonaws.services.quicksight.model.RegisteredUserDashboardEmbeddingConfiguration;
import
com.amazonaws.services.quicksight.model.AnonymousUserEmbeddingExperienceConfiguration;
import
com.amazonaws.services.quicksight.model.GenerateEmbedUrlForAnonymousUserRequest;
import
com.amazonaws.services.quicksight.model.GenerateEmbedUrlForAnonymousUserResult;
import com.amazonaws.services.quicksight.model.SessionTag;
```

```

/**
 * Class to call QuickSight AWS SDK to generate embed url for anonymous user.
 */
public class GenerateEmbedUrlForAnonymousUserExample {

    private final AmazonQuickSight quickSightClient;

    public GenerateEmbedUrlForAnonymousUserExample() {
        quickSightClient = AmazonQuickSightClientBuilder
            .standard()
            .withRegion(Regions.US_EAST_1.getName())
            .withCredentials(new AWSCredentialsProvider() {
                @Override
                public AWSCredentials getCredentials() {
                    // provide actual IAM access key and secret key here
                    return new BasicAWSCredentials("access-key", "secret-key");
                }

                @Override
                public void refresh() {
                }
            })
            .build();
    }

    public String GenerateEmbedUrlForAnonymousUser(
        final String accountId, // YOUR AWS ACCOUNT ID
        final String initialDashboardId, // DASHBOARD ID TO WHICH THE CONSTRUCTED
        URL POINTS.
        final String namespace, // ANONYMOUS EMBEDDING REQUIRES SPECIFYING A VALID
        NAMESPACE FOR WHICH YOU WANT THE EMBEDDING URL
        final List<String> authorizedResourceArns, // DASHBOARD ARN LIST TO EMBED
        final List<String> allowedDomains, // RUNTIME ALLOWED DOMAINS FOR EMBEDDING
        final List<SessionTag> sessionTags // SESSION TAGS USED FOR ROW-LEVEL
        SECURITY
    ) throws Exception {
        AnonymousUserEmbeddingExperienceConfiguration experienceConfiguration = new
        AnonymousUserEmbeddingExperienceConfiguration();
        AnonymousUserDashboardEmbeddingConfiguration dashboardConfiguration = new
        AnonymousUserDashboardEmbeddingConfiguration();
        dashboardConfiguration.setInitialDashboardId(initialDashboardId);
        experienceConfiguration.setDashboard(dashboardConfiguration);
    }
}

```

```

        GenerateEmbedUrlForAnonymousUserRequest
generateEmbedUrlForAnonymousUserRequest = new
GenerateEmbedUrlForAnonymousUserRequest()
    .withAwsAccountId(accountId)
    .withNamespace(namespace)
    .withAuthorizedResourceArns(authorizedResourceArns)
    .withExperienceConfiguration(experienceConfiguration)
    .withSessionTags(sessionTags)
    .withSessionLifetimeInMinutes(600L); // OPTIONAL: VALUE CAN BE
[15-600]. DEFAULT: 600
    .withAllowedDomains(allowedDomains);

        GenerateEmbedUrlForAnonymousUserResult dashboardEmbedUrl =
quickSightClient.generateEmbedUrlForAnonymousUser(generateEmbedUrlForAnonymousUserRequest);

        return dashboardEmbedUrl.getEmbedUrl();
    }

}

```

JavaScript

```

global.fetch = require('node-fetch');
const AWS = require('aws-sdk');

function generateEmbedUrlForAnonymousUser(
accountId, // YOUR AWS ACCOUNT ID
initialDashboardId, // DASHBOARD ID TO WHICH THE CONSTRUCTED URL POINTS
quicksightNamespace, // VALID NAMESPACE WHERE YOU WANT TO DO NOAUTH EMBEDDING
authorizedResourceArns, // DASHBOARD ARN LIST TO EMBED
allowedDomains, // RUNTIME ALLOWED DOMAINS FOR EMBEDDING
sessionTags, // SESSION TAGS USED FOR ROW-LEVEL SECURITY
generateEmbedUrlForAnonymousUserCallback, // GENERATEEMBEDURLFORANONYMOUSUSER SUCCESS
CALLBACK METHOD
errorCallback // GENERATEEMBEDURLFORANONYMOUSUSER ERROR CALLBACK METHOD
) {
const experienceConfiguration = {
    "DashboardVisual": {
        "InitialDashboardVisualId": {
            "DashboardId": "dashboard_id",
            "SheetId": "sheet_id",
            "VisualId": "visual_id"
        }
    }
}

```

```
    }  
};  
  
const generateEmbedUrlForAnonymousUserParams = {  
  "AwsAccountId": accountId,  
  "Namespace": quicksightNamespace,  
  "AuthorizedResourceArns": authorizedResourceArns,  
  "AllowedDomains": allowedDomains,  
  "ExperienceConfiguration": experienceConfiguration,  
  "SessionTags": sessionTags,  
  "SessionLifetimeInMinutes": 600  
};  
  
const quicksightClient = new AWS.QuickSight({  
  region: process.env.AWS_REGION,  
  credentials: {  
    accessKeyId: AccessKeyId,  
    secretAccessKey: SecretAccessKey,  
    sessionToken: SessionToken,  
    expiration: Expiration  
  }  
});  
  
quicksightClient.generateEmbedUrlForAnonymousUser(generateEmbedUrlForAnonymousUserParams,  
function(err, data) {  
  if (err) {  
    console.log(err, err.stack);  
    errorCallback(err);  
  } else {  
    const result = {  
      "statusCode": 200,  
      "headers": {  
        "Access-Control-Allow-Origin": "*", // USE YOUR WEBSITE DOMAIN TO  
SECURE ACCESS TO THIS API  
        "Access-Control-Allow-Headers": "Content-Type"  
      },  
      "body": JSON.stringify(data),  
      "isBase64Encoded": false  
    }  
    generateEmbedUrlForAnonymousUserCallback(result);  
  }  
});  
}
```


Python3

```
import json
import boto3
from botocore.exceptions import ClientError
import time

# Create QuickSight and STS clients
quicksightClient = boto3.client('quicksight', region_name='us-west-2')
sts = boto3.client('sts')

# Function to generate embedded URL for anonymous user
# accountId: YOUR AWS ACCOUNT ID
# quicksightNamespace: VALID NAMESPACE WHERE YOU WANT TO DO NOAUTH EMBEDDING
# authorizedResourceArns: DASHBOARD ARN LIST TO EMBED
# allowedDomains: RUNTIME ALLOWED DOMAINS FOR EMBEDDING
# dashboardId: DASHBOARD ID TO WHICH THE CONSTRUCTED URL POINTS
# sessionTags: SESSION TAGS USED FOR ROW-LEVEL SECURITY
def generateEmbedUrlForAnonymousUser(accountId, quicksightNamespace,
    authorizedResourceArns, allowedDomains, dashboardId, sessionTags):
    try:
        response = quicksightClient.generate_embed_url_for_anonymous_user(
            AwsAccountId = accountId,
            Namespace = quicksightNamespace,
            AuthorizedResourceArns = authorizedResourceArns,
            AllowedDomains = allowedDomains,
            ExperienceConfiguration = {
                "Dashboard": {
                    "InitialDashboardId": dashboardId
                }
            },
            SessionTags = sessionTags,
            SessionLifetimeInMinutes = 600
        )

        return {
            'statusCode': 200,
            'headers': {"Access-Control-Allow-Origin": "*", "Access-Control-Allow-Headers":
"Content-Type"},
            'body': json.dumps(response),
            'isBase64Encoded': bool('false')
        }
    except ClientError as e:
        print(e)
```

```
return "Error generating embeddedURL: " + str(e)
```

Node.js

The following example shows the JavaScript (Node.js) that you can use on the app server to generate the URL for the embedded dashboard. You can use this URL in your website or app to display the dashboard.

Example

```
const AWS = require('aws-sdk');
const https = require('https');

var quicksightClient = new AWS.Service({
  apiConfig: require('./quicksight-2018-04-01.min.json'),
  region: 'us-east-1',
});

quicksightClient.generateEmbedUrlForAnonymousUser({
  'AwsAccountId': '111122223333',
  'Namespace' : 'default',
  'AuthorizedResourceArns': authorizedResourceArns,
  'AllowedDomains': allowedDomains,
  'ExperienceConfiguration': experienceConfiguration,
  'SessionTags': sessionTags,
  'SessionLifetimeInMinutes': 600

}, function(err, data) {
  console.log('Errors: ');
  console.log(err);
  console.log('Response: ');
  console.log(data);
});
```

Example

```
//The URL returned is over 900 characters. For this example, we've shortened the string
for
  //readability and added ellipsis to indicate that it's incomplete.
  {
    Status: 200,
    EmbedUrl: 'https://quicksightdomain/embed/12345/dashboards/67890..',
    RequestId: '7bee030e-f191-45c4-97fe-d9faf0e03713'
```

```
}
```

.NET/C#

The following example shows the .NET/C# code that you can use on the app server to generate the URL for the embedded dashboard. You can use this URL in your website or app to display the dashboard.

Example

```
using System;
using Amazon.QuickSight;
using Amazon.QuickSight.Model;

var quicksightClient = new AmazonQuickSightClient(
    AccessKey,
    SecretAccessKey,
    sessionToken,
    Amazon.RegionEndpoint.USEast1);

try
{
    Console.WriteLine(
        quicksightClient.GenerateEmbedUrlForAnonymousUserAsync(new
GenerateEmbedUrlForAnonymousUserRequest
    {
        AwsAccountId = "111122223333",
        Namespace = default,
        AuthorizedResourceArns = authorizedResourceArns,
        AllowedDomains = allowedDomains,
        ExperienceConfiguration = experienceConfiguration,
        SessionTags = sessionTags,
        SessionLifetimeInMinutes = 600,
    }).Result.EmbedUrl
    );
} catch (Exception ex) {
    Console.WriteLine(ex.Message);
}
```

AWS CLI

To assume the role, choose one of the following AWS Security Token Service (AWS STS) API operations:

- [AssumeRole](#) – Use this operation when you're using an IAM identity to assume the role.
- [AssumeRoleWithWebIdentity](#) – Use this operation when you're using a web identity provider to authenticate your user.
- [AssumeRoleWithSaml](#) – Use this operation when you're using Security Assertion Markup Language (SAML) to authenticate your users.

The following example shows the CLI command to set the IAM role. The role needs to have permissions enabled for `quicksight:GenerateEmbedUrlForAnonymousUser`.

```
aws sts assume-role \  
  --role-arn "arn:aws:iam::11112222333:role/QuickSightEmbeddingAnonymousPolicy" \  
  --role-session-name anonymous caller
```

The `assume-role` operation returns three output parameters: the access key, the secret key, and the session token.

Note

If you get an `ExpiredToken` error when calling the `AssumeRole` operation, this is probably because the previous `SESSION_TOKEN` is still in the environment variables. Clear this by setting the following variables:

- `AWS_ACCESS_KEY_ID`
- `AWS_SECRET_ACCESS_KEY`
- `AWS_SESSION_TOKEN`

The following example shows how to set these three parameters in the CLI. If you're using a Microsoft Windows machine, use `set` instead of `export`.

```
export AWS_ACCESS_KEY_ID      = "access_key_from_assume_role"  
export AWS_SECRET_ACCESS_KEY = "secret_key_from_assume_role"  
export AWS_SESSION_TOKEN     = "session_token_from_assume_role"
```

Running these commands sets the role session ID of the user visiting your website to `embedding_quicksight_dashboard_role/QuickSightEmbeddingAnonymousPolicy`. The role session ID is made up of the role name from `role-arn` and the `role-session-name` value.

Using the unique role session ID for each user ensures that appropriate permissions are set for each visiting user. It also keeps each session separate and distinct. If you're using an array of web servers, for example for load balancing, and a session is reconnected to a different server, a new session begins.

To get a signed URL for the dashboard, call `generate-embed-url-for-anonymous-user` from the app server. This returns the embeddable dashboard URL. The following example shows how to generate the URL for an embedded dashboard using a server-side call for users who are making anonymous visits to your web portal or app.

```
aws quicksight generate-embed-url-for-anonymous-user \
--aws-account-id 111122223333 \
--namespace default-or-something-else \
--session-lifetime-in-minutes 15 \
--authorized-resource-arns '["dashboard-arn-1","dashboard-arn-2"]' \
--allowed-domains '["domain1","domain2"]' \
--session-tags '["Key": tag-key-1,"Value": tag-value-1,{"Key": tag-key-1,"Value": tag-
value-1}]' \
--experience-configuration
'DashboardVisual={InitialDashboardVisualId={DashboardId=dashboard_id,SheetId=sheet_id,VisualId=
```

For more information about using this operation, see [GenerateEmbedUrlForAnonymousUser](#). You can use this and other API operations in your own code.

Step 3: Embed the dashboard URL

Applies to: Enterprise Edition

Intended audience: Amazon Quick developers

In the following section, you can find out how you can use the [Amazon Quick Sight Embedding SDK](#) (JavaScript) to embed the dashboard URL from step 2 in your website or application page. With the SDK, you can do the following:

- Place the dashboard on an HTML page.
- Pass parameters into the dashboard.
- Handle error states with messages that are customized to your application.

Call the `GenerateEmbedUrlForAnonymousUser` API operation to generate the URL that you can embed in your app. This URL is valid for 5 minutes, and the resulting session is valid for 10 hours. The API operation provides the URL with an `auth_code` that enables a single-sign on session.

The following shows an example response from `generate-embed-url-for-anonymous-user`.

```
//The URL returned is over 900 characters. For this example, we've shortened the string
//for
//readability and added ellipsis to indicate that it's incomplete.
{
  "Status": "200",
  "EmbedUrl": "https://quicksightdomain/embed/12345/dashboards/67890..",
  "RequestId": "7bee030e-f191-45c4-97fe-d9faf0e03713"
}
```

Embed this dashboard in your web page by using the [Amazon Quick Sight Embedding SDK](#) or by adding this URL into an `iframe`. If you set a fixed height and width number (in pixels), Amazon Quick Sight uses those and doesn't change your visual as your window resizes. If you set a relative percent height and width, Amazon Quick Sight provides a responsive layout that is modified as your window size changes. By using the Amazon Quick Sight Embedding SDK, you can also control parameters within the dashboard and receive callbacks in terms of page load completion and errors.

The domain that is going to host embedded dashboards must be on the *allow list*, the list of approved domains for your Quick subscription. This requirement protects your data by keeping unapproved domains from hosting embedded dashboards. For more information about adding domains for embedded dashboards, see [Allow listing domains at runtime with the Amazon Quick Sight API](#).

The following example shows how to use the generated URL. This code resides on your app server.

SDK 2.0

```
<!DOCTYPE html>
<html>

<head>
  <title>Dashboard Embedding Example</title>
  <script src="https://unpkg.com/amazon-quicksight-embedding-sdk@2.0.0/dist/
quicksight-embedding-js-sdk.min.js"></script>
  <script type="text/javascript">
```

```

const embedDashboard = async() => {
  const {
    createEmbeddingContext,
  } = QuickSightEmbedding;

  const embeddingContext = await createEmbeddingContext({
    onChange: (changeEvent, metadata) => {
      console.log('Context received a change', changeEvent, metadata);
    },
  });

  const frameOptions = {
    url: '<YOUR_EMBED_URL>',
    container: '#experience-container',
    height: "700px",
    width: "1000px",
    onChange: (changeEvent, metadata) => {
      switch (changeEvent.eventName) {
        case 'FRAME_MOUNTED': {
          console.log("Do something when the experience frame is
mounted.");
          break;
        }
        case 'FRAME_LOADED': {
          console.log("Do something when the experience frame is
loaded.");
          break;
        }
      }
    },
  };

  const contentOptions = {
    parameters: [
      {
        Name: 'country',
        Values: [
          'United States'
        ],
      },
      {
        Name: 'states',
        Values: [
          'California',

```

```

        'Washington'
    ]
}
],
locale: "en-US",
sheetOptions: {
    initialSheetId: '<YOUR_SHEETID>',
    singleSheet: false,
    emitSizeChangedEventOnSheetChange: false,
},
toolbarOptions: {
    export: false,
    undoRedo: false,
    reset: false
},
attributionOptions: {
    overlayContent: false,
},
onMessage: async (messageEvent, experienceMetadata) => {
    switch (messageEvent.eventName) {
        case 'CONTENT_LOADED': {
            console.log("All visuals are loaded. The title of the
document:", messageEvent.message.title);
            break;
        }
        case 'ERROR_OCCURRED': {
            console.log("Error occurred while rendering the experience.
Error code:", messageEvent.message.errorCode);
            break;
        }
        case 'PARAMETERS_CHANGED': {
            console.log("Parameters changed. Changed parameters:",
messageEvent.message.changedParameters);
            break;
        }
        case 'SELECTED_SHEET_CHANGED': {
            console.log("Selected sheet changed. Selected sheet:",
messageEvent.message.selectedSheet);
            break;
        }
        case 'SIZE_CHANGED': {
            console.log("Size changed. New dimensions:",
messageEvent.message);
            break;
    }
}
}

```



```

        }
        case 'MODAL_OPENED': {
            window.scrollTo({
                top: 0 // iframe top position
            });
            break;
        }
    },
    },
    };
    const embeddedDashboardExperience = await
embeddingContext.embedDashboard(frameOptions, contentOptions);

    const selectCountryElement = document.getElementById('country');
    selectCountryElement.addEventListener('change', (event) => {
        embeddedDashboardExperience.setParameters([
            {
                Name: 'country',
                Values: event.target.value
            }
        ]);
    });
});
</script>
</head>

<body onload="embedDashboard()">
    <span>
        <label for="country">Country</label>
        <select id="country" name="country">
            <option value="United States">United States</option>
            <option value="Mexico">Mexico</option>
            <option value="Canada">Canada</option>
        </select>
    </span>
    <div id="experience-container"></div>
</body>

</html>

```

SDK 1.0

```
<!DOCTYPE html>
```

```
<html>

<head>
  <title>Basic Embed</title>
  <script src="https://unpkg.com/amazon-quicksight-embedding-sdk@1.0.15/dist/
quicksight-embedding-js-sdk.min.js"></script>
  <script type="text/javascript">
    var dashboard
    function onDashboardLoad(payload) {
      console.log("Do something when the dashboard is fully loaded.");
    }

    function onError(payload) {
      console.log("Do something when the dashboard fails loading");
    }

    function embedDashboard() {
      var containerDiv = document.getElementById("embeddingContainer");
      var options = {
        // replace this dummy url with the one generated via embedding API
        url: "https://us-east-1.quicksight.aws.amazon.com/sn/dashboards/
dashboardId?isauthcode=true&identityprovider=quicksight&code=authcode",
        container: containerDiv,
        parameters: {
          country: "United States"
        },
        scrolling: "no",
        height: "700px",
        width: "1000px",
        locale: "en-US",
        footerPaddingEnabled: true
      };
      dashboard = QuickSightEmbedding.embedDashboard(options);
      dashboard.on("error", onError);
      dashboard.on("load", onDashboardLoad);
    }

    function onCountryChange(obj) {
      dashboard.setParameters({country: obj.value});
    }
  </script>
</head>

<body onload="embedDashboard()">
```

```
<span>
  <label for="country">Country</label>
  <select id="country" name="country" onchange="onCountryChange(this)">
    <option value="United States">United States</option>
    <option value="Mexico">Mexico</option>
    <option value="Canada">Canada</option>
  </select>
</span>
<div id="embeddingContainer"></div>
</body>

</html>
```

For this example to work, make sure to use the Amazon Quick Sight Embedding SDK to load the embedded dashboard on your website using JavaScript. To get your copy, do one of the following:

- Download the [Amazon Quick Sight Embedding SDK](#) from GitHub. This repository is maintained by a group of Amazon Quick Sight developers.
- Download the latest Amazon Quick Sight Embedding SDK version from <https://www.npmjs.com/package/amazon-quicksight-embedding-sdk>.
- If you use npm for JavaScript dependencies, download and install it by running the following command.

```
npm install amazon-quicksight-embedding-sdk
```

Enabling executive summaries in embedded dashboards

Applies to: Enterprise Edition

Intended audience: Amazon Quick developers

You can enable executive summaries in your embedded dashboards. When enabled, registered users can generate executive summaries that provide a summary of all insights that Amazon Quick Sight has generated for the dashboard. Executive summaries make it easier for readers to find key insights and information about a dashboard. For more information about how users generate an

executive summary of a dashboard, see [Generate an executive summary of an Amazon Quick Sight dashboard](#).

Note

Executive summaries are only available in embedded dashboards for registered users, and cannot be enabled in embedded dashboards for anonymous or unregistered users.

To enable executive summaries in embedded dashboards for registered users

- Follow the steps in [Embedding Amazon Quick Sight dashboards for registered users](#) to embed a dashboard with the following changes:
 - When generating the URL in Step 2, set `Enabled: true` in the `ExecutiveSummary` parameter in the [GenerateEmbedUrlForRegisteredUser](#) or [GenerateEmbedUrlForRegisteredUserWithIdentity](#) as shown in the following example:

```
ExperienceConfiguration: {
  Dashboard: {
    InitialDashboardId: dashboard_id,
    FeatureConfigurations: {
      AmazonQInQuickSight: {
        ExecutiveSummary: {
          Enabled: true
        }
      }
    }
  }
}
```

- When embedding the dashboard URL with the Amazon Quick Sight Embedding SDK in Step 3, set `executiveSummary: true` in `contentOptions`, as shown in the following example:

```
const contentOptions = {
  toolbarOptions: {
    executiveSummary: true
  }
};
```

Embedding Amazon Quick Sight visuals with the Amazon Quick Sight APIs

You can embed individual visuals that are a part of a published dashboard in your application with the Amazon Quick Sight API.

Topics

- [Embedding Amazon Quick Sight visuals for registered users](#)
- [Embedding Amazon Quick Sight visuals for anonymous \(unregistered\) users](#)

Embedding Amazon Quick Sight visuals for registered users

Applies to: Enterprise Edition

Intended audience: Amazon Quick developers

In the following sections, you can find detailed information about how to set up embedded Amazon Quick Sight visuals for registered users of Amazon Quick Sight.

Topics

- [Step 1: Set up permissions](#)
- [Step 2: Generate the URL with the authentication code attached](#)
- [Step 3: Embed the visual URL](#)

Step 1: Set up permissions

In the following section, you can find out how to set up permissions for the backend application or web server. This task requires administrative access to IAM.

Each user who accesses a visual assumes a role that gives them Amazon Quick Sight access and permissions to the visual. To make this possible, create an IAM role in your AWS account. Associate an IAM policy with the role to provide permissions to any user who assumes it. The IAM role needs to provide permissions to retrieve embedding URLs for a specific user pool. With the help of the wildcard character *, you can grant the permissions to generate a URL for all users in a specific namespace, or for a subset of users in specific namespaces. For this, you add `quicksight:GenerateEmbedUrlForRegisteredUser`.

You can create a condition in your IAM policy that limits the domains that developers can list in the `AllowedDomains` parameter of a `GenerateEmbedUrlForAnonymousUser` API operation. The `AllowedDomains` parameter is an optional parameter. It grants you as a developer the option to override the static domains that are configured in the **Manage Amazon Quick Sight** menu. Instead, you can list up to three domains or subdomains that can access a generated URL. This URL is then embedded in the website that you create. Only the domains that are listed in the parameter can access the embedded dashboard. Without this condition, you can list any domain on the internet in the `AllowedDomains` parameter.

To limit the domains that developers can use with this parameter, add an `AllowedEmbeddingDomains` condition to your IAM policy. For more information about the `AllowedDomains` parameter, see [GenerateEmbedUrlForRegisteredUser](#) in the *Amazon Quick Sight API Reference*.

Security best practice for IAM condition operators

Improperly configured IAM condition operators can allow unauthorized access to your embedded Quick resources through URL variations. When using the `quicksight:AllowedEmbeddingDomains` condition key in your IAM policies, use condition operators that either allow specific domains or deny all domains that are not specifically allowed. For more information about IAM condition operators, see [IAM JSON policy elements: Condition operators](#) in the IAM User Guide.

Many different URL variations can point to the same resource. For example, the following URLs all resolve to the same content:

- `https://example.com`
- `https://example.com/`
- `https://Example.com`

If your policy uses operators that do not account for these URL variations, an attacker can bypass your restrictions by providing equivalent URL variations.

You must validate that your IAM policy uses appropriate condition operators to prevent bypass vulnerabilities and ensure that only your intended domains can access your embedded resources.

The following sample policy provides these permissions.

Additionally, if you are creating first-time users who will be Amazon Quick Sight readers, make sure to add the `quicksight:RegisterUser` permission in the policy.

The following sample policy provides permission to retrieve an embedding URL for first-time users who are to be Amazon Quick Sight readers.

Finally, your application's IAM identity must have a trust policy associated with it to allow access to the role that you just created. This means that when a user accesses your application, your application can assume the role on the user's behalf and provision the user in Amazon Quick Sight. The following example shows a sample trust policy.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowLambdaFunctionsToAssumeThisRole",
      "Effect": "Allow",
      "Principal": {
        "Service": "lambda.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    },
    {
      "Sid": "AllowEC2InstancesToAssumeThisRole",
      "Effect": "Allow",
      "Principal": {
        "Service": "ec2.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

For more information regarding trust policies for OpenID Connect or SAML authentication, see the following sections of the *IAM User Guide*:

- [Creating a Role for Web Identity or OpenID Connect Federation \(Console\)](#)
- [Creating a Role for SAML 2.0 Federation \(Console\)](#)

Step 2: Generate the URL with the authentication code attached

In the following section, you can find out how to authenticate your Amazon Quick Sight user and get the embeddable visual URL on your application server. If you plan to embed visuals for IAM or Amazon Quick Sight identity types, share the visual with the Amazon Quick Sight users.

When a Amazon Quick Sight user accesses your app, the app assumes the IAM role on the Amazon Quick Sight user's behalf. Then it adds the user to Amazon Quick Sight, if that Amazon Quick Sight user doesn't already exist. Next, it passes an identifier as the unique role session ID.

Performing the described steps ensures that each viewer of the visual is uniquely provisioned in Amazon Quick Sight. It also enforces per-user settings, such as the row-level security and dynamic defaults for parameters.

The following examples perform the IAM authentication on the Amazon Quick Sight user's behalf. This code runs on your app server.

Java

```
import com.amazonaws.auth.AWSCredentials;
import com.amazonaws.auth.BasicAWSCredentials;
import com.amazonaws.auth.AWSCredentialsProvider;
import com.amazonaws.regions.Regions;
import com.amazonaws.services.quicksight.AmazonQuickSight;
import com.amazonaws.services.quicksight.AmazonQuickSightClientBuilder;
import com.amazonaws.services.quicksight.model.DashboardVisualId;
import
    com.amazonaws.services.quicksight.model.GenerateEmbedUrlForRegisteredUserRequest;
import com.amazonaws.services.quicksight.model.GenerateEmbedUrlForRegisteredUserResult;
import
    com.amazonaws.services.quicksight.model.RegisteredUserDashboardVisualEmbeddingConfiguration;
import
    com.amazonaws.services.quicksight.model.RegisteredUserEmbeddingExperienceConfiguration;

import java.util.List;

/**
 * Class to call QuickSight AWS SDK to get url for Visual embedding.
 */
public class GenerateEmbedUrlForRegisteredUserTest {

    private final AmazonQuickSight quickSightClient;
```



```

public GenerateEmbedUrlForRegisteredUserTest() {
    this.quickSightClient = AmazonQuickSightClientBuilder
        .standard()
        .withRegion(Regions.US_EAST_1.getName())
        .withCredentials(new AWSCredentialsProvider() {
            @Override
            public AWSCredentials getCredentials() {
                // provide actual IAM access key and secret key here
                return new BasicAWSCredentials("access-key", "secret-key");
            }

            @Override
            public void refresh() {
            }
        })
        .build();
}

public String getEmbedUrl(
    final String accountId, // AWS Account ID
    final String dashboardId, // Dashboard ID of the dashboard to embed
    final String sheetId, // Sheet ID of the sheet to embed
    final String visualId, // Visual ID of the visual to embed
    final List<String> allowedDomains, // Runtime allowed domains for embedding
    final String userArn // Registered user arn of the user that you want to
    provide embedded visual. Refer to Get Embed Url section in developer portal to find
    out how to get user arn for a QuickSight user.
) throws Exception {
    final DashboardVisualId dashboardVisual = new DashboardVisualId()
        .withDashboardId(dashboardId)
        .withSheetId(sheetId)
        .withVisualId(visualId);

    final RegisteredUserDashboardVisualEmbeddingConfiguration
registeredUserDashboardVisualEmbeddingConfiguration
        = new RegisteredUserDashboardVisualEmbeddingConfiguration()
            .withInitialDashboardVisualId(dashboardVisual);

    final RegisteredUserEmbeddingExperienceConfiguration
registeredUserEmbeddingExperienceConfiguration
        = new RegisteredUserEmbeddingExperienceConfiguration()

.withDashboardVisual(registeredUserDashboardVisualEmbeddingConfiguration);
    final GenerateEmbedUrlForRegisteredUserRequest
generateEmbedUrlForRegisteredUserRequest

```

```

        = new GenerateEmbedUrlForRegisteredUserRequest()
            .withAwsAccountId(accountId)
            .withUserArn(userArn)

        .withExperienceConfiguration(registeredUserEmbeddingExperienceConfiguration)
            .withAllowedDomains(allowedDomains);

        final GenerateEmbedUrlForRegisteredUserResult
generateEmbedUrlForRegisteredUserResult =
quickSightClient.generateEmbedUrlForRegisteredUser(generateEmbedUrlForRegisteredUserRequest);

        return generateEmbedUrlForRegisteredUserResult.getEmbedUrl();
    }
}

```

JavaScript

```

global.fetch = require('node-fetch');
const AWS = require('aws-sdk');

function generateEmbedUrlForRegisteredUser(
    accountId, // Your AWS account ID
    dashboardId, // Dashboard ID to which the constructed URL points
    sheetId, // Sheet ID to which the constructed URL points
    visualId, // Visual ID to which the constructed URL points
    openIdToken, // Cognito-based token
    userArn, // registered user arn
    roleArn, // IAM user role to use for embedding
    sessionName, // Session name for the roleArn assume role
    allowedDomains, // Runtime allowed domain for embedding
    getEmbedUrlCallback, // GetEmbedUrl success callback method
    errorCallback // GetEmbedUrl error callback method
) {
    const stsClient = new AWS.STS();
    let stsParams = {
        RoleSessionName: sessionName,
        WebIdentityToken: openIdToken,
        RoleArn: roleArn
    };

    stsClient.assumeRoleWithWebIdentity(stsParams, function(err, data) {
        if (err) {
            console.log('Error assuming role');
        }
    });
}

```

```

        console.log(err, err.stack);
        errorCallback(err);
    } else {
        const getDashboardParams = {
            "AwsAccountId": accountId,
            "ExperienceConfiguration": {
                "DashboardVisual": {
                    "InitialDashboardVisualId": {
                        "DashboardId": dashboardId,
                        "SheetId": sheetId,
                        "VisualId": visualId
                    }
                }
            },
            "UserArn": userArn,
            "AllowedDomains": allowedDomains,
            "SessionLifetimeInMinutes": 600
        };

        const quicksightGetDashboard = new AWS.QuickSight({
            region: process.env.AWS_REGION,
            credentials: {
                accessKeyId: data.Credentials.AccessKeyId,
                secretAccessKey: data.Credentials.SecretAccessKey,
                sessionToken: data.Credentials.SessionToken,
                expiration: data.Credentials.Expiration
            }
        });

        quicksightGetDashboard.generateEmbedUrlForRegisteredUser(getDashboardParams,
        function(err, data) {
            if (err) {
                console.log(err, err.stack);
                errorCallback(err);
            } else {
                const result = {
                    "statusCode": 200,
                    "headers": {
                        "Access-Control-Allow-Origin": "*", // Use your website
                        "Access-Control-Allow-Headers": "Content-Type"
                    },
                    "body": JSON.stringify(data),

```

```

        "isBase64Encoded": false
    }
    getEmbedUrlCallback(result);
}
});
}
});
}
}

```

Python3

```

import json
import boto3
from botocore.exceptions import ClientError

sts = boto3.client('sts')

# Function to generate embedded URL
# accountId: AWS account ID
# dashboardId: Dashboard ID to embed
# sheetId: SHEET ID to embed from the dashboard
# visualId: Id for the Visual you want to embedded from the dashboard sheet.
# userArn: arn of registered user
# allowedDomains: Runtime allowed domain for embedding
# roleArn: IAM user role to use for embedding
# sessionName: session name for the roleArn assume role
def getEmbeddingURL(accountId, dashboardId, sheetId, visualId, userArn, allowedDomains,
    roleArn, sessionName):
    try:
        assumedRole = sts.assume_role(
            RoleArn = roleArn,
            RoleSessionName = sessionName,
        )
    except ClientError as e:
        return "Error assuming role: " + str(e)
    else:
        assumedRoleSession = boto3.Session(
            aws_access_key_id = assumedRole['Credentials']['AccessKeyId'],
            aws_secret_access_key = assumedRole['Credentials']['SecretAccessKey'],
            aws_session_token = assumedRole['Credentials']['SessionToken'],
        )
        try:

```

```

    quicksightClient = assumedRoleSession.client('quicksight', region_name='us-
west-2')

    response = quicksightClient.generate_embed_url_for_registered_user(
        AwsAccountId=accountId,
        ExperienceConfiguration = {
            'DashboardVisual': {
                'InitialDashboardVisualId': {
                    'DashboardId': dashboardId,
                    'SheetId': sheetId,
                    'VisualId': visualId
                }
            },
        },
        UserArn = userArn,
        AllowedDomains = allowedDomains,
        SessionLifetimeInMinutes = 600
    )

    return {
        'statusCode': 200,
        'headers': {"Access-Control-Allow-Origin": "*", "Access-Control-Allow-
Headers": "Content-Type"},
        'body': json.dumps(response),
        'isBase64Encoded': bool('false')
    }
except ClientError as e:
    return "Error generating embedding url: " + str(e)

```

Node.js

The following example shows the JavaScript (Node.js) that you can use on the app server to generate the URL for the embedded dashboard. You can use this URL in your website or app to display the dashboard.

Example

```

const AWS = require('aws-sdk');
const https = require('https');

var quicksightClient = new AWS.Service({
    apiConfig: require('./quicksight-2018-04-01.min.json'),
    region: 'us-east-1',
});

```

```
quicksightClient.generateEmbedUrlForRegisteredUser({
  'AwsAccountId': '111122223333',
  'ExperienceConfiguration': {
    'DashboardVisual': {
      'InitialDashboardVisualId': {
        'DashboardId': 'dashboard_id',
        'SheetId': 'sheet_id',
        'VisualId': 'visual_id'
      }
    }
  },
  'UserArn': 'REGISTERED_USER_ARN',
  'AllowedDomains': allowedDomains,
  'SessionLifetimeInMinutes': 100
}, function(err, data) {
  console.log('Errors: ');
  console.log(err);
  console.log('Response: ');
  console.log(data);
});
```

Example

```
//The URL returned is over 900 characters. For this example, we've shortened the string
//for
//readability and added ellipsis to indicate that it's incomplete.
{
  "Status": "200",
  "EmbedUrl": "https://quicksightdomain/embed/12345/dashboards/67890/
sheets/12345/visuals/67890...",
  "RequestId": "7bee030e-f191-45c4-97fe-d9faf0e03713"
}
```

.NET/C#

The following example shows the .NET/C# code that you can use on the app server to generate the URL for the embedded dashboard. You can use this URL in your website or app to display the dashboard.

Example

```
using System;
```

```

using Amazon.QuickSight;
using Amazon.QuickSight.Model;

namespace GenerateDashboardEmbedUrlForRegisteredUser
{
    class Program
    {
        static void Main(string[] args)
        {
            var quicksightClient = new AmazonQuickSightClient(
                AccessKey,
                SecretAccessKey,
                SessionToken,
                Amazon.RegionEndpoint.USEast1);

            try
            {
                DashboardVisualId dashboardVisual = new DashboardVisualId
                {
                    DashboardId = "dashboard_id",
                    SheetId = "sheet_id",
                    VisualId = "visual_id"
                };

                RegisteredUserDashboardVisualEmbeddingConfiguration
registeredUserDashboardVisualEmbeddingConfiguration
                    = new RegisteredUserDashboardVisualEmbeddingConfiguration
                    {
                        InitialDashboardVisualId = dashboardVisual
                    };

                RegisteredUserEmbeddingExperienceConfiguration
registeredUserEmbeddingExperienceConfiguration
                    = new RegisteredUserEmbeddingExperienceConfiguration
                    {
                        DashboardVisual =
registeredUserDashboardVisualEmbeddingConfiguration
                    };

                Console.WriteLine(
                    quicksightClient.GenerateEmbedUrlForRegisteredUserAsync(new
GenerateEmbedUrlForRegisteredUserRequest
                    {
                        AwsAccountId = "111122223333",

```

```

        ExperienceConfiguration =
registeredUserEmbeddingExperienceConfiguration,
        UserArn = "REGISTERED_USER_ARN",
        AllowedDomains = allowedDomains,
        SessionLifetimeInMinutes = 100
    })).Result.EmbedUrl
    );
} catch (Exception ex) {
    Console.WriteLine(ex.Message);
}
}
}
}

```

AWS CLI

To assume the role, choose one of the following AWS Security Token Service (AWS STS) API operations:

- [AssumeRole](#) – Use this operation when you're using an IAM identity to assume the role.
- [AssumeRoleWithWebIdentity](#) – Use this operation when you're using a web identity provider to authenticate your user.
- [AssumeRoleWithSaml](#) – Use this operation when you're using SAML to authenticate your users.

The following example shows the CLI command to set the IAM role. The role needs to have permissions enabled for `quicksight:GenerateEmbedUrlForRegisteredUser`. If you are taking a just-in-time approach to add users when they first open a dashboard, the role also needs permissions enabled for `quicksight:RegisterUser`.

```

aws sts assume-role \
  --role-arn "arn:aws:iam::111122223333:role/embedding_quicksight_visual_role" \
  --role-session-name john.doe@example.com

```

The `assume-role` operation returns three output parameters: the access key, the secret key, and the session token.

Note

If you get an `ExpiredToken` error when calling the `AssumeRole` operation, this is probably because the previous `SESSION_TOKEN` is still in the environment variables. Clear this by setting the following variables:

- `AWS_ACCESS_KEY_ID`
- `AWS_SECRET_ACCESS_KEY`
- `AWS_SESSION_TOKEN`

The following example shows how to set these three parameters in the CLI. If you're using a Microsoft Windows machine, use `set` instead of `export`.

```
export AWS_ACCESS_KEY_ID      = "access_key_from_assume_role"
export AWS_SECRET_ACCESS_KEY  = "secret_key_from_assume_role"
export AWS_SESSION_TOKEN      = "session_token_from_assume_role"
```

Running these commands sets the role session ID of the user visiting your website to `embedding_quicksight_visual_role/john.doe@example.com`. The role session ID is made up of the role name from `role-arn` and the `role-session-name` value. Using the unique role session ID for each user ensures that appropriate permissions are set for each user. It also prevents any throttling of user access. *Throttling* is a security feature that prevents the same user from accessing Amazon Quick Sight from multiple locations.

The role session ID also becomes the user name in Amazon Quick Sight. You can use this pattern to provision your users in Amazon Quick Sight ahead of time, or to provision them the first time they access the dashboard.

The following example shows the CLI command that you can use to provision a user. For more information about [RegisterUser](#), [DescribeUser](#), and other Amazon Quick Sight API operations, see the [Amazon Quick Sight API Reference](#).

```
aws quicksight register-user \
  --aws-account-id 111122223333 \
  --namespace default \
  --identity-type IAM \
  --iam-arn "arn:aws:iam::111122223333:role/embedding_quicksight_visual_role" \
  --user-role READER \
```

```
--user-name jhnd \
--session-name "john.doe@example.com" \
--email john.doe@example.com \
--region us-east-1 \
--custom-permissions-name TeamA1
```

If the user is authenticated through Microsoft AD, you don't need to use `RegisterUser` to set them up. Instead, they should be automatically subscribed the first time they access Amazon Quick Sight. For Microsoft AD users, you can use `DescribeUser` to get the user ARN.

The first time a user accesses Amazon Quick Sight, you can also add this user to the group that the visual is shared with. The following example shows the CLI command to add a user to a group.

```
aws quicksight create-group-membership \
  --aws-account-id=111122223333 \
  --namespace=default \
  --group-name=financeusers \
  --member-name="embedding_quicksight_visual_role/john.doe@example.com"
```

You now have a user of your app who is also a user of Amazon Quick Sight, and who has access to the visual.

Finally, to get a signed URL for the visual, call `generate-embed-url-for-registered-user` from the app server. This returns the embeddable visual URL. The following example shows how to generate the URL for an embedded visual using a server-side call for users authenticated through AWS Managed Microsoft AD or single sign-on (IAM Identity Center).

```
aws quicksight generate-embed-url-for-registered-user \
  --aws-account-id 111122223333 \
  --session-lifetime-in-minutes 600 \
  --user-arn arn:aws:quicksight:us-east-1:111122223333:user/default/
embedding_quicksight_visual_role/embeddingsession \
  --allowed-domains '["domain1", "domain2"]' \
  --experience-configuration
'DashboardVisual={InitialDashboardVisualId={DashboardId=dashboard_id, SheetId=sheet_id, VisualId=
```

For more information about using this operation, see [GenerateEmbedUrlForRegisteredUser](#). You can use this and other API operations in your own code.

Step 3: Embed the visual URL

In the following section, you can find out how you can use the [Amazon Quick Sight Embedding SDK](#) (JavaScript) to embed the visual URL from step 3 in your website or application page. With the SDK, you can do the following:

- Place the visual on an HTML page.
- Pass parameters into the visual.
- Handle error states with messages that are customized to your application.

Call the `GenerateEmbedUrlForRegisteredUser` API operation to generate the URL that you can embed in your app. This URL is valid for 5 minutes, and the resulting session is valid for up to 10 hours. The API operation provides the URL with an `auth_code` that enables a single-sign on session.

The following shows an example response from `generate-embed-url-for-registered-user`. The *quicksightdomain* in this example is the URL that you use to access your Amazon Quick Sight account.

```
//The URL returned is over 900 characters. For this example, we've shortened the string
//for
//readability and added ellipsis to indicate that it's incomplete.
{
  "Status": "200",
  "EmbedUrl": "https://quicksightdomain/embed/12345/dashboards/67890/
sheets/12345/visuals/67890...",
  "RequestId": "7bee030e-f191-45c4-97fe-d9faf0e03713"
}
```

Embed this visual in your webpage by using the [Amazon Quick Sight Embedding SDK](#) or by adding this URL into an `iframe`. If you set a fixed height and width number (in pixels), Amazon Quick Sight uses those and doesn't change your visual as your window resizes. If you set a relative percent height and width, Amazon Quick Sight provides a responsive layout that is modified as your window size changes. By using the Amazon Quick Sight Embedding SDK, you can also control parameters within the visual and receive callbacks in terms of page load completion and errors.

The domain that is going to host embedded visuals and dashboards must be on the *allow list*, the list of approved domains for your Quick subscription. This requirement protects your data

by keeping unapproved domains from hosting embedded visuals and dashboards. For more information about adding domains for embedded visuals and dashboards, see [Allow listing domains at runtime with the Amazon Quick Sight API](#).

The following example shows how to use the generated URL. This code is generated on your app server.

SDK 2.0

```
<!DOCTYPE html>
<html>

  <head>
    <title>Visual Embedding Example</title>
    <script src="https://unpkg.com/amazon-quicksight-embedding-sdk@2.0.0/dist/
quicksight-embedding-js-sdk.min.js"></script>
    <script type="text/javascript">
      const embedVisual = async() => {
        const {
          createEmbeddingContext,
        } = QuickSightEmbedding;

        const embeddingContext = await createEmbeddingContext({
          onChange: (changeEvent, metadata) => {
            console.log('Context received a change', changeEvent,
metadata);
          },
        });

        const frameOptions = {
          url: "<YOUR_EMBED_URL>", // replace this value with the url
generated via embedding API
          container: '#experience-container',
          height: "700px",
          width: "1000px",
          onChange: (changeEvent, metadata) => {
            switch (changeEvent.eventName) {
              case 'FRAME_MOUNTED': {
                console.log("Do something when the experience frame is
mounted.");
                break;
              }
              case 'FRAME_LOADED': {
```

```

        console.log("Do something when the experience frame is
loaded.");
        break;
    }
    },
};

const contentOptions = {
    parameters: [
        {
            Name: 'country',
            Values: ['United States'],
        },
        {
            Name: 'states',
            Values: [
                'California',
                'Washington'
            ]
        }
    ],
    locale: "en-US",
    onMessage: async (messageEvent, experienceMetadata) => {
        switch (messageEvent.eventName) {
            case 'CONTENT_LOADED': {
                console.log("All visuals are loaded. The title of the
document:", messageEvent.message.title);
                break;
            }
            case 'ERROR_OCCURRED': {
                console.log("Error occurred while rendering the
experience. Error code:", messageEvent.message.errorCode);
                break;
            }
            case 'PARAMETERS_CHANGED': {
                console.log("Parameters changed. Changed parameters:",
messageEvent.message.changedParameters);
                break;
            }
            case 'SIZE_CHANGED': {
                console.log("Size changed. New dimensions:",
messageEvent.message);
                break;
            }
        }
    }
};

```

```

        }
      },
    },
  ];
  const embeddedVisualExperience = await
embeddingContext.embedVisual(frameOptions, contentOptions);

  const selectCountryElement = document.getElementById('country');
  selectCountryElement.addEventListener('change', (event) => {
    embeddedVisualExperience.setParameters([
      {
        Name: 'country',
        Values: event.target.value
      }
    ]);
  });
});
</script>
</head>

<body onload="embedVisual()">
  <span>
    <label for="country">Country</label>
    <select id="country" name="country">
      <option value="United States">United States</option>
      <option value="Mexico">Mexico</option>
      <option value="Canada">Canada</option>
    </select>
  </span>
  <div id="experience-container"></div>
</body>

</html>

```

SDK 1.0

```

<!DOCTYPE html>
<html>

  <head>
    <title>Visual Embedding Example</title>
    <!-- You can download the latest QuickSight embedding SDK version from https://
www.npmjs.com/package/amazon-quicksight-embedding-sdk -->

```

```

    <!-- Or you can do "npm install amazon-quicksight-embedding-sdk", if you use
    npm for javascript dependencies -->
    <script src="./quicksight-embedding-js-sdk.min.js"></script>
    <script type="text/javascript">
        let embeddedVisualExperience;
        function onVisualLoad(payload) {
            console.log("Do something when the visual is fully loaded.");
        }

        function onError(payload) {
            console.log("Do something when the visual fails loading");
        }

        function embedVisual() {
            const containerDiv = document.getElementById("embeddingContainer");
            const options = {
                url: "<YOUR_EMBED_URL>", // replace this value with the url
generated via embedding API
                container: containerDiv,
                parameters: {
                    country: "United States"
                },
                height: "700px",
                width: "1000px",
                locale: "en-US"
            };
            embeddedVisualExperience = QuickSightEmbedding.embedVisual(options);
            embeddedVisualExperience.on("error", onError);
            embeddedVisualExperience.on("load", onVisualLoad);
        }

        function onCountryChange(obj) {
            embeddedVisualExperience.setParameters({country: obj.value});
        }
    </script>
</head>

<body onload="embedVisual()">
    <span>
        <label for="country">Country</label>
        <select id="country" name="country" onchange="onCountryChange(this)">
            <option value="United States">United States</option>
            <option value="Mexico">Mexico</option>
            <option value="Canada">Canada</option>

```

```
        </select>
    </span>
    <div id="embeddingContainer"></div>
</body>

</html>
```

For this example to work, make sure to use the Amazon Quick Sight Embedding SDK to load the embedded visual on your website using JavaScript. To get your copy, do one of the following:

- Download the [Amazon Quick Sight Embedding SDK](#) from GitHub. This repository is maintained by a group of Amazon Quick Sight developers.
- Download the latest embedding SDK version from <https://www.npmjs.com/package/amazon-quicksight-embedding-sdk>.
- If you use npm for JavaScript dependencies, download and install it by running the following command.

```
npm install amazon-quicksight-embedding-sdk
```

Embedding Amazon Quick Sight visuals for anonymous (unregistered) users

Applies to: Enterprise Edition

Intended audience: Amazon Quick developers

In the following sections, you can find detailed information about how to set up embedded Amazon Quick Sight visuals for anonymous (unregistered) users.

Topics

- [Step 1: Set up permissions](#)
- [Step 2: Generate the URL with the authentication code attached](#)
- [Step 3: Embed the visual URL](#)

Step 1: Set up permissions

Applies to: Enterprise Edition

Intended audience: Amazon Quick developers

In the following section, you can find out how to set up permissions for the backend application or web server. This task requires administrative access to IAM.

Each user who accesses a visual assumes a role that gives them Amazon Quick Sight access and permissions to the visual. To make this possible, create an IAM role in your AWS account. Associate an IAM policy with the role to provide permissions to any user who assumes it.

You can create a condition in your IAM policy that limits the domains that developers can list in the `AllowedDomains` parameter of a `GenerateEmbedUrlForAnonymousUser` API operation. The `AllowedDomains` parameter is an optional parameter. It grants you as a developer the option to override the static domains that are configured in the **Manage Amazon Quick Sight** menu. Instead, you can list up to three domains or subdomains that can access a generated URL. This URL is then embedded in the website that you create. Only the domains that are listed in the parameter can access the embedded dashboard. Without this condition, you can list any domain on the internet in the `AllowedDomains` parameter.

To limit the domains that developers can use with this parameter, add an `AllowedEmbeddingDomains` condition to your IAM policy. For more information about the `AllowedDomains` parameter, see [GenerateEmbedUrlForAnonymousUser](#) in the *Amazon Quick Sight API Reference*.

Security best practice for IAM condition operators

Improperly configured IAM condition operators can allow unauthorized access to your embedded Quick resources through URL variations. When using the `quicksight:AllowedEmbeddingDomains` condition key in your IAM policies, use condition operators that either allow specific domains or deny all domains that are not specifically allowed. For more information about IAM condition operators, see [IAM JSON policy elements: Condition operators](#) in the IAM User Guide.

Many different URL variations can point to the same resource. For example, the following URLs all resolve to the same content:

- `https://example.com`
- `https://example.com/`
- `https://Example.com`

If your policy uses operators that do not account for these URL variations, an attacker can bypass your restrictions by providing equivalent URL variations.

You must validate that your IAM policy uses appropriate condition operators to prevent bypass vulnerabilities and ensure that only your intended domains can access your embedded resources.

Your application's IAM identity must have a trust policy associated with it to allow access to the role that you just created. This means that when a user accesses your application, your application can assume the role on the user's behalf to open the visual. The following example shows a sample trust policy.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowLambdaFunctionsToAssumeThisRole",
      "Effect": "Allow",
      "Principal": {
        "Service": "lambda.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    },
    {
      "Sid": "AllowEC2InstancesToAssumeThisRole",
      "Effect": "Allow",
      "Principal": {
        "Service": "ec2.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

```
}  
]  
}
```

For more information regarding trust policies, see [Temporary security credentials in IAM](#) in the *IAM User Guide*.

Step 2: Generate the URL with the authentication code attached

Applies to: Enterprise Edition

Intended audience: Amazon Quick developers

In the following section, you can find how to authenticate on behalf of the anonymous visitor and get the embeddable visual URL on your application server.

When a user accesses your app, the app assumes the IAM role on the user's behalf. Then it adds the user to Amazon Quick Sight, if that user doesn't already exist. Next, it passes an identifier as the unique role session ID.

The following examples perform the IAM authentication on the user's behalf. It passes an identifier as the unique role session ID. This code runs on your app server.

Java

```
import com.amazonaws.auth.AWSCredentials;  
import com.amazonaws.auth.BasicAWSCredentials;  
import com.amazonaws.auth.AWSCredentialsProvider;  
import com.amazonaws.regions.Regions;  
import com.amazonaws.services.quicksight.AmazonQuickSight;  
import com.amazonaws.services.quicksight.AmazonQuickSightClientBuilder;  
import  
    com.amazonaws.services.quicksight.model.AnonymousUserDashboardVisualEmbeddingConfiguration;  
import  
    com.amazonaws.services.quicksight.model.AnonymousUserEmbeddingExperienceConfiguration;  
import com.amazonaws.services.quicksight.model.DashboardVisualId;  
import com.amazonaws.services.quicksight.model.GenerateEmbedUrlForAnonymousUserRequest;  
import com.amazonaws.services.quicksight.model.GenerateEmbedUrlForAnonymousUserResult;
```

```

import com.amazonaws.services.quicksight.model.SessionTag;

import java.util.List;

/**
 * Class to call QuickSight AWS SDK to get url for Visual embedding.
 */
public class GenerateEmbedUrlForAnonymousUserTest {
    private final AmazonQuickSight quickSightClient;

    public GenerateEmbedUrlForAnonymousUserTest() {
        this.quickSightClient = AmazonQuickSightClientBuilder
            .standard()
            .withRegion(Regions.US_EAST_1.getName())
            .withCredentials(new AWSCredentialsProvider() {
                @Override
                public AWSCredentials getCredentials() {
                    // provide actual IAM access key and secret key here
                    return new BasicAWSCredentials("access-key", "secret-key");
                }

                @Override
                public void refresh() {
                }
            })
            .build();
    }

    public String getEmbedUrl(
        final String accountId, // AWS Account ID
        final String namespace, // Anonymous embedding required specifying a valid
namespace for which you want the embedding URL
        final List<String> authorizedResourceArns, // Dashboard arn list of
dashboard visuals to embed
        final String dashboardId, // Dashboard ID of the dashboard to embed
        final String sheetId, // Sheet ID of the sheet to embed
        final String visualId, // Visual ID of the visual to embed
        final List<String> allowedDomains, // Runtime allowed domains for embedding
        final List<SessionTag> sessionTags // Session tags used for row-level
security
    ) throws Exception {
        final DashboardVisualId dashboardVisual = new DashboardVisualId()
            .withDashboardId(dashboardId)

```

```

        .withSheetId(sheetId)
        .withVisualId(visualId);
    final AnonymousUserDashboardVisualEmbeddingConfiguration
anonymousUserDashboardVisualEmbeddingConfiguration
        = new AnonymousUserDashboardVisualEmbeddingConfiguration()
            .withInitialDashboardVisualId(dashboardVisual);
    final AnonymousUserEmbeddingExperienceConfiguration
anonymousUserEmbeddingExperienceConfiguration
        = new AnonymousUserEmbeddingExperienceConfiguration()

.withDashboardVisual(anonymousUserDashboardVisualEmbeddingConfiguration);
    final GenerateEmbedUrlForAnonymousUserRequest
generateEmbedUrlForAnonymousUserRequest
        = new GenerateEmbedUrlForAnonymousUserRequest()
            .withAwsAccountId(accountId)
            .withNamespace(namespace)
            // authorizedResourceArns should contain ARN of dashboard used below in
ExperienceConfiguration
            .withAuthorizedResourceArns(authorizedResourceArns)

.withExperienceConfiguration(anonymousUserEmbeddingExperienceConfiguration)
            .withAllowedDomains(allowedDomains)
            .withSessionTags(sessionTags)
            .withSessionLifetimeInMinutes(600L);

    final GenerateEmbedUrlForAnonymousUserResult
generateEmbedUrlForAnonymousUserResult
        =
quickSightClient.generateEmbedUrlForAnonymousUser(generateEmbedUrlForAnonymousUserRequest);

    return generateEmbedUrlForAnonymousUserResult.getEmbedUrl();
}
}

```

JavaScript

```

global.fetch = require('node-fetch');
const AWS = require('aws-sdk');

function generateEmbedUrlForAnonymousUser(
    accountId, // Your AWS account ID
    dashboardId, // Dashboard ID to which the constructed url points
    sheetId, // Sheet ID to which the constructed url points

```

```

visualId, // Visual ID to which the constructed url points
quicksightNamespace, // valid namespace where you want to do embedding
authorizedResourceArns, // dashboard arn list of dashboard visuals to embed
allowedDomains, // runtime allowed domains for embedding
sessionTags, // session tags used for row-level security
generateEmbedUrlForAnonymousUserCallback, // success callback method
errorCallback // error callback method
) {
const experienceConfiguration = {
  "DashboardVisual": {
    "InitialDashboardVisualId": {
      "DashboardId": dashboardId,
      "SheetId": sheetId,
      "VisualId": visualId
    }
  }
};

const generateEmbedUrlForAnonymousUserParams = {
  "AwsAccountId": accountId,
  "Namespace": quicksightNamespace,
  // authorizedResourceArns should contain ARN of dashboard used below in
ExperienceConfiguration
  "AuthorizedResourceArns": authorizedResourceArns,
  "AllowedDomains": allowedDomains,
  "ExperienceConfiguration": experienceConfiguration,
  "SessionTags": sessionTags,
  "SessionLifetimeInMinutes": 600
};

const quicksightClient = new AWS.QuickSight({
  region: process.env.AWS_REGION,
  credentials: {
    accessKeyId: AccessKeyId,
    secretAccessKey: SecretAccessKey,
    sessionToken: SessionToken,
    expiration: Expiration
  }
});

quicksightClient.generateEmbedUrlForAnonymousUser(generateEmbedUrlForAnonymousUserParams,
function(err, data) {
  if (err) {

```

```

        console.log(err, err.stack);
        errorCallback(err);
    } else {
        const result = {
            "statusCode": 200,
            "headers": {
                "Access-Control-Allow-Origin": "*", // USE YOUR WEBSITE DOMAIN TO
SECURE ACCESS TO THIS API
                "Access-Control-Allow-Headers": "Content-Type"
            },
            "body": JSON.stringify(data),
            "isBase64Encoded": false
        }
        generateEmbedUrlForAnonymousUserCallback(result);
    }
});
}

```

Python3

```

import json
import boto3
from botocore.exceptions import ClientError
import time

# Create QuickSight and STS clients
quicksightClient = boto3.client('quicksight', region_name='us-west-2')
sts = boto3.client('sts')

# Function to generate embedded URL for anonymous user
# accountId: YOUR AWS ACCOUNT ID
# quicksightNamespace: VALID NAMESPACE WHERE YOU WANT TO DO NOAUTH EMBEDDING
# authorizedResourceArns: DASHBOARD ARN LIST TO EMBED
# allowedDomains: RUNTIME ALLOWED DOMAINS FOR EMBEDDING
# experienceConfiguration: DASHBOARD ID, SHEET ID and VISUAL ID TO WHICH THE
CONSTRUCTED URL POINTS
# Example experienceConfig -> 'DashboardVisual': {
#     'InitialDashboardVisualId': {
#         'DashboardId': 'dashboardId',
#         'SheetId': 'sheetId',
#         'VisualId': 'visualId'
#     }
# },

```

```
# sessionTags: SESSION TAGS USED FOR ROW-LEVEL SECURITY
def generateEmbedUrlForAnonymousUser(accountId, quicksightNamespace,
    authorizedResourceArns, allowedDomains, experienceConfiguration, sessionTags):
    try:
        response = quicksightClient.generate_embed_url_for_anonymous_user(
            AwsAccountId = accountId,
            Namespace = quicksightNamespace,
            AuthorizedResourceArns = authorizedResourceArns,
            AllowedDomains = allowedDomains,
            ExperienceConfiguration = experienceConfiguration,
            SessionTags = sessionTags,
            SessionLifetimeInMinutes = 600
        )

        return {
            'statusCode': 200,
            'headers': {"Access-Control-Allow-Origin": "*", "Access-Control-Allow-Headers": "Content-Type"},
            'body': json.dumps(response),
            'isBase64Encoded': bool('false')
        }
    except ClientError as e:
        print(e)
        return "Error generating embeddedURL: " + str(e)
```

Node.js

The following example shows the JavaScript (Node.js) that you can use on the app server to generate the URL for the embedded dashboard. You can use this URL in your website or app to display the dashboard.

Example

```
const AWS = require('aws-sdk');
const https = require('https');

var quicksightClient = new AWS.Service({
    apiConfig: require('./quicksight-2018-04-01.min.json'),
    region: 'us-east-1',
});

quicksightClient.generateEmbedUrlForAnonymousUser({
    'AwsAccountId': '111122223333',
```



```

    'Namespace' : 'default',
    // authorizedResourceArns should contain ARN of dashboard used below in
    ExperienceConfiguration
    'AuthorizedResourceArns': authorizedResourceArns,
    'ExperienceConfiguration': {
        'DashboardVisual': {
            'InitialDashboardVisualId': {
                'DashboardId': 'dashboard_id',
                'SheetId': 'sheet_id',
                'VisualId': 'visual_id'
            }
        }
    },
    'AllowedDomains': allowedDomains,
    'SessionTags': sessionTags,
    'SessionLifetimeInMinutes': 600
}, function(err, data) {
    console.log('Errors: ');
    console.log(err);
    console.log('Response: ');
    console.log(data);
});

```

Example

```

//The URL returned is over 900 characters. For this example, we've shortened the string
for
//readability and added ellipsis to indicate that it's incomplete.
{
    "Status": "200",
    "EmbedUrl": "https://quicksightdomain/embed/12345/dashboards/67890/
sheets/12345/visuals/67890...",
    "RequestId": "7bee030e-f191-45c4-97fe-d9faf0e03713"
}

```

.NET/C#

The following example shows the .NET/C# code that you can use on the app server to generate the URL for the embedded dashboard. You can use this URL in your website or app to display the dashboard.

Example

```
using System;
using Amazon.QuickSight;
using Amazon.QuickSight.Model;

namespace GenerateDashboardEmbedUrlForAnonymousUser
{
    class Program
    {
        static void Main(string[] args)
        {
            var quicksightClient = new AmazonQuickSightClient(
                AccessKey,
                SecretAccessKey,
                SessionToken,
                Amazon.RegionEndpoint.USEast1);
            try
            {
                DashboardVisualId dashboardVisual = new DashboardVisualId
                {
                    DashboardId = "dashboard_id",
                    SheetId = "sheet_id",
                    VisualId = "visual_id"
                };

                AnonymousUserDashboardVisualEmbeddingConfiguration
anonymousUserDashboardVisualEmbeddingConfiguration
                    = new AnonymousUserDashboardVisualEmbeddingConfiguration
                    {
                        InitialDashboardVisualId = dashboardVisual
                    };

                AnonymousUserEmbeddingExperienceConfiguration
anonymousUserEmbeddingExperienceConfiguration
                    = new AnonymousUserEmbeddingExperienceConfiguration
                    {
                        DashboardVisual =
anonymousUserDashboardVisualEmbeddingConfiguration
                    };

                Console.WriteLine(
```

```

        quicksightClient.GenerateEmbedUrlForAnonymousUserAsync(new
GenerateEmbedUrlForAnonymousUserRequest
        {
            AwsAccountId = "111222333444",
            Namespace = default,
            // authorizedResourceArns should contain ARN of dashboard used
below in ExperienceConfiguration
            AuthorizedResourceArns = { "dashboard_id" },
            ExperienceConfiguration =
anonymousUserEmbeddingExperienceConfiguration,
            SessionTags = sessionTags,
            SessionLifetimeInMinutes = 600,
        }).Result.EmbedUrl
    );
} catch (Exception ex) {
    Console.WriteLine(ex.Message);
}
}
}
}
}

```

AWS CLI

To assume the role, choose one of the following AWS Security Token Service (AWS STS) API operations:

- [AssumeRole](#) – Use this operation when you're using an IAM identity to assume the role.
- [AssumeRoleWithWebIdentity](#) – Use this operation when you're using a web identity provider to authenticate your user.
- [AssumeRoleWithSaml](#) – Use this operation when you're using Security Assertion Markup Language (SAML) to authenticate your users.

The following example shows the CLI command to set the IAM role. The role needs to have permissions enabled for `quicksight:GenerateEmbedUrlForAnonymousUser`.

```

aws sts assume-role \
    --role-arn "arn:aws:iam::11112222333:role/QuickSightEmbeddingAnonymousPolicy" \
    --role-session-name anonymous caller

```

The `assume-role` operation returns three output parameters: the access key, the secret key, and the session token.

Note

If you get an `ExpiredToken` error when calling the `AssumeRole` operation, this is probably because the previous `SESSION_TOKEN` is still in the environment variables. Clear this by setting the following variables:

- `AWS_ACCESS_KEY_ID`
- `AWS_SECRET_ACCESS_KEY`
- `AWS_SESSION_TOKEN`

The following example shows how to set these three parameters in the CLI. If you're using a Microsoft Windows machine, use `set` instead of `export`.

```
export AWS_ACCESS_KEY_ID      = "access_key_from_assume_role"  
export AWS_SECRET_ACCESS_KEY  = "secret_key_from_assume_role"  
export AWS_SESSION_TOKEN      = "session_token_from_assume_role"
```

Running these commands sets the role session ID of the user visiting your website to `embedding_quicksight_visual_role/QuickSightEmbeddingAnonymousPolicy`. The role session ID is made up of the role name from `role-arn` and the `role-session-name` value. Using the unique role session ID for each user ensures that appropriate permissions are set for each visiting user. It also keeps each session separate and distinct. If you're using an array of web servers, for example for load balancing, and a session is reconnected to a different server, a new session begins.

To get a signed URL for the visual, call `generate-embed-url-for-anonymous-user` from the app server. This returns the embeddable visual URL. The following example shows how to generate the URL for an embedded visual using a server-side call for users who are making anonymous visits to your web portal or app.

```
aws quicksight generate-embed-url-for-anonymous-user \  
  --aws-account-id 111122223333 \  
  --namespace default-or-something-else \  
  --session-lifetime-in-minutes 15 \  
  --authorized-resource-arns '["dashboard-arn-1","dashboard-arn-2"]' \  
  \
```

```
--allowed-domains '["domain1", "domain2"]' \
--session-tags '["Key": tag-key-1, "Value": tag-value-1, {"Key": tag-
key-1, "Value": tag-value-1}]' \
--experience-configuration
'DashboardVisual={InitialDashboardVisualId={DashboardId=dashboard_id, SheetId=sheet_id, VisualId=
```

For more information about using this operation, see [GenerateEmbedUrlForAnonymousUser](#). You can use this and other API operations in your own code.

Step 3: Embed the visual URL

Applies to: Enterprise Edition

Intended audience: Amazon Quick developers

In the following section, you can find out how you can use the [Amazon Quick Sight Embedding SDK](#) (JavaScript) to embed the visual URL from step 2 in your website or application page. With the SDK, you can do the following:

- Place the visual on an HTML page.
- Pass parameters into the visual.
- Handle error states with messages that are customized to your application.

Call the `GenerateEmbedUrlForAnonymousUser` API operation to generate the URL that you can embed in your app. This URL is valid for 5 minutes, and the resulting session is valid for 10 hours. The API operation provides the URL with an authorization (auth) code that enables a single-sign on session.

The following shows an example response from `generate-embed-url-for-anonymous-user`. The *quicksightdomain* in this example is the URL that you use to access your Amazon Quick Sight account.

```
//The URL returned is over 900 characters. For this example, we've shortened the string
for
//readability and added ellipsis to indicate that it's incomplete.
{
    "Status": "200",
```

```
"EmbedUrl": "https://quicksightdomain/embed/12345/dashboards/67890/sheets/12345/visuals/67890...",
"RequestId": "7bee030e-f191-45c4-97fe-d9faf0e03713"
}
```

Embed this visual in your web page by using the Amazon Quick Sight [Embedding SDK](#) or by adding this URL into an iframe. If you set a fixed height and width number (in pixels), Amazon Quick Sight uses those and doesn't change your visual as your window resizes. If you set a relative percent height and width, Amazon Quick Sight provides a responsive layout that is modified as your window size changes. By using the Amazon Quick Sight Embedding SDK, you can also control parameters within the visual and receive callbacks in terms of visual load completion and errors.

The domain that is going to host embedded visual must be on the *allow list*, the list of approved domains for your Quick subscription. This requirement protects your data by keeping unapproved domains from hosting embedded visuals and dashboards. For more information about adding domains for embedded visuals and dashboards, see [Allow listing domains at runtime with the Amazon Quick Sight API](#).

The following example shows how to use the generated URL. This code resides on your app server.

SDK 2.0

```
<!DOCTYPE html>
<html>

  <head>
    <title>Visual Embedding Example</title>
    <script src="https://unpkg.com/amazon-quicksight-embedding-sdk@2.0.0/dist/quicksight-embedding-js-sdk.min.js"></script>
    <script type="text/javascript">
      const embedVisual = async() => {
        const {
          createEmbeddingContext,
        } = QuickSightEmbedding;

        const embeddingContext = await createEmbeddingContext({
          onChange: (changeEvent, metadata) => {
            console.log('Context received a change', changeEvent,
metadata);
          },
        });
      };
```

```

        const frameOptions = {
            url: "<YOUR_EMBED_URL>", // replace this value with the url
generated via embedding API
            container: '#experience-container',
            height: "700px",
            width: "1000px",
            onChange: (changeEvent, metadata) => {
                switch (changeEvent.eventName) {
                    case 'FRAME_MOUNTED': {
                        console.log("Do something when the experience frame is
mounted.");

                        break;
                    }
                    case 'FRAME_LOADED': {
                        console.log("Do something when the experience frame is
loaded.");

                        break;
                    }
                }
            },
        };

        const contentOptions = {
            parameters: [
                {
                    Name: 'country',
                    Values: ['United States'],
                },
                {
                    Name: 'states',
                    Values: [
                        'California',
                        'Washington'
                    ]
                }
            ],
            locale: "en-US",
            onMessage: async (messageEvent, experienceMetadata) => {
                switch (messageEvent.eventName) {
                    case 'CONTENT_LOADED': {
                        console.log("All visuals are loaded. The title of the
document:", messageEvent.message.title);
                        break;
                    }
                }
            }
        }

```

```

        case 'ERROR_OCCURRED': {
            console.log("Error occurred while rendering the
experience. Error code:", messageEvent.message.errorCode);
            break;
        }
        case 'PARAMETERS_CHANGED': {
            console.log("Parameters changed. Changed parameters:",
messageEvent.message.changedParameters);
            break;
        }
        case 'SIZE_CHANGED': {
            console.log("Size changed. New dimensions:",
messageEvent.message);
            break;
        }
    },
    },
    };
    const embeddedVisualExperience = await
embeddingContext.embedVisual(frameOptions, contentOptions);

    const selectCountryElement = document.getElementById('country');
    selectCountryElement.addEventListener('change', (event) => {
        embeddedVisualExperience.setParameters([
            {
                Name: 'country',
                Values: event.target.value
            }
        ]);
    });
});
</script>
</head>

<body onload="embedVisual()">
    <span>
        <label for="country">Country</label>
        <select id="country" name="country">
            <option value="United States">United States</option>
            <option value="Mexico">Mexico</option>
            <option value="Canada">Canada</option>
        </select>
    </span>
    <div id="experience-container"></div>

```



```
</body>

</html>
```

SDK 1.0

```
<!DOCTYPE html>
<html>

  <head>
    <title>Visual Embedding Example</title>
    <!-- You can download the latest QuickSight embedding SDK version from https://
www.npmjs.com/package/amazon-quicksight-embedding-sdk -->
    <!-- Or you can do "npm install amazon-quicksight-embedding-sdk", if you use
npm for javascript dependencies -->
    <script src="./quicksight-embedding-js-sdk.min.js"></script>
    <script type="text/javascript">
      let embeddedVisualExperience;
      function onVisualLoad(payload) {
        console.log("Do something when the visual is fully loaded.");
      }

      function onError(payload) {
        console.log("Do something when the visual fails loading");
      }

      function embedVisual() {
        const containerDiv = document.getElementById("embeddingContainer");
        const options = {
          url: "<YOUR_EMBED_URL>", // replace this value with the url
generated via embedding API
          container: containerDiv,
          parameters: {
            country: "United States"
          },
          height: "700px",
          width: "1000px",
          locale: "en-US"
        };
        embeddedVisualExperience = QuickSightEmbedding.embedVisual(options);
        embeddedVisualExperience.on("error", onError);
        embeddedVisualExperience.on("load", onVisualLoad);
      }
    </script>
  </head>
</html>
```

```
        function onCountryChange(obj) {
            embeddedVisualExperience.setParameters({country: obj.value});
        }
    </script>
</head>

<body onload="embedVisual()">
    <span>
        <label for="country">Country</label>
        <select id="country" name="country" onchange="onCountryChange(this)">
            <option value="United States">United States</option>
            <option value="Mexico">Mexico</option>
            <option value="Canada">Canada</option>
        </select>
    </span>
    <div id="embeddingContainer"></div>
</body>

</html>
```

For this example to work, make sure to use the Amazon Quick Sight Embedding SDK to load the embedded visual on your website using JavaScript. To get your copy, do one of the following:

- Download the [Amazon Quick Sight Embedding SDK](#) from GitHub. This repository is maintained by a group of Amazon Quick Sight developers.
- Download the latest QuickSight embedding SDK version from <https://www.npmjs.com/package/amazon-quicksight-embedding-sdk>.
- If you use npm for JavaScript dependencies, download and install it by running the following command.

```
npm install amazon-quicksight-embedding-sdk
```

Embedding the full functionality of the Amazon Quick Sight console for registered users

Important

Amazon Quick Sight has new API operations for embedding analytics:
GenerateEmbedUrlForAnonymousUser and GenerateEmbedUrlForRegisteredUser.

You can still use the `GetDashboardEmbedUrl` and `GetSessionEmbedUrl` API operations to embed dashboards and the Amazon Quick Sight console, but they don't contain the latest embedding capabilities. For more information about embedding using the old API operations, see [Embedding analytics using the `GetDashboardEmbedURL` and `GetSessionEmbedURL` API operations](#).

Applies to: Enterprise Edition

Intended audience: Amazon Quick developers

With Enterprise edition, in addition to providing read-only dashboards you can also provide the Amazon Quick Sight console experience in a custom-branded authoring portal. Using this approach, you allow your users to create data sources, datasets, and analyses. In the same interface, they can create, publish, and view dashboards. If you want to restrict some of those permissions, you can also do that.

Users who access Amazon Quick Sight through an embedded console need to belong to the author or admin security cohort. Readers don't have enough access to use the Amazon Quick Sight console for authoring, regardless of whether it's embedded or part of the AWS Management Console. However, authors and admins can still access embedded dashboards. If you want to restrict permissions to some of the authoring features, you can add a custom permissions profile to the user with the [UpdateUser](#) API operation. Use the [RegisterUser](#) API operation to add a new user with a custom permission profile attached. For more information, see the following sections:

- For information about creating custom roles by defining custom console permissions, see [Customizing Access to the Amazon Quick Sight Console](#).
- For information about using namespaces to isolate multitenancy users, groups, and Amazon Quick Sight assets, see [Amazon Quick Sight Namespaces](#).
- For information about adding your own branding to an embedded Amazon Quick Sight console, see [Using Themes in Amazon Quick Sight](#) and the [QuickSight Theme API Operations](#).

In the following sections, you can find detailed information about how to set up embedded Amazon Quick Sight dashboards for registered users.

Topics

- [Step 1: Set up permissions](#)
- [Step 2: Generate the URL with the authentication code attached](#)
- [Step 3: Embed the console session URL](#)
- [Enabling Generative BI features in embedded consoles for registered users](#)

Step 1: Set up permissions

In the following section, you can find out how to set up permissions for the backend application or web server. This task requires administrative access to IAM.

Each user who accesses a Amazon Quick Sight assumes a role that gives them Amazon Quick Sight access and permissions to the console session. To make this possible, create an IAM role in your AWS account. Associate an IAM policy with the role to provide permissions to any user who assumes it. Add `quicksight:RegisterUser` permissions to ensure that the reader can access Amazon Quick Sight in a read-only fashion, and not have access to any other data or creation capability. The IAM role also needs to provide permissions to retrieve console session URLs. For this, you add `quicksight:GenerateEmbedUrlForRegisteredUser`.

You can create a condition in your IAM policy that limits the domains that developers can list in the `AllowedDomains` parameter of a `GenerateEmbedUrlForAnonymousUser` API operation. The `AllowedDomains` parameter is an optional parameter. It grants you as a developer the option to override the static domains that are configured in the **Manage Amazon Quick Sight** menu. Instead, you can list up to three domains or subdomains that can access a generated URL. This URL is then embedded in the website that you create. Only the domains that are listed in the parameter can access the embedded dashboard. Without this condition, you can list any domain on the internet in the `AllowedDomains` parameter.

Security best practice for IAM condition operators

Improperly configured IAM condition operators can allow unauthorized access to your embedded Quick resources through URL variations. When using the `quicksight:AllowedEmbeddingDomains` condition key in your IAM policies, use condition operators that either allow specific domains or deny all domains that are not specifically allowed. For more information about IAM condition operators, see [IAM JSON policy elements: Condition operators](#) in the IAM User Guide.

Many different URL variations can point to the same resource. For example, the following URLs all resolve to the same content:

- `https://example.com`
- `https://example.com/`
- `https://Example.com`

If your policy uses operators that do not account for these URL variations, an attacker can bypass your restrictions by providing equivalent URL variations.

You must validate that your IAM policy uses appropriate condition operators to prevent bypass vulnerabilities and ensure that only your intended domains can access your embedded resources.

The following sample policy provides these permissions.

The following sample policy provides permission to retrieve a console session URL. You can use the policy without `quicksight:RegisterUser` if you are creating users before they access an embedded session.

Finally, your application's IAM identity must have a trust policy associated with it to allow access to the role that you just created. This means that when a user accesses your application, your application can assume the role on the user's behalf and provision the user in Amazon Quick Sight. The following example shows a sample trust policy.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowLambdaFunctionsToAssumeThisRole",
      "Effect": "Allow",
      "Principal": {
        "Service": "lambda.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    },
    {
```

```
        "Sid": "AllowEC2InstancesToAssumeThisRole",
        "Effect": "Allow",
        "Principal": {
            "Service": "ec2.amazonaws.com"
        },
        "Action": "sts:AssumeRole"
    }
]
```

For more information regarding trust policies for OpenID Connect or SAML authentication, see the following sections of the *IAM User Guide*:

- [Creating a Role for Web Identity or OpenID Connect Federation \(Console\)](#)
- [Creating a Role for SAML 2.0 Federation \(Console\)](#)

Step 2: Generate the URL with the authentication code attached

In the following section, you can find out how to authenticate your user and get the embeddable console session URL on your application server.

When a user accesses your app, the app assumes the IAM role on the user's behalf. Then it adds the user to Amazon Quick Sight, if that user doesn't already exist. Next, it passes an identifier as the unique role session ID.

Performing the described steps ensures that each viewer of the console session is uniquely provisioned in Amazon Quick Sight. It also enforces per-user settings, such as the row-level security and dynamic defaults for parameters.

The following examples perform the IAM authentication on the user's behalf. This code runs on your app server.

Java

```
import com.amazonaws.auth.AWSCredentials;
import com.amazonaws.auth.BasicAWSCredentials;
import com.amazonaws.auth.AWSCredentialsProvider;
import com.amazonaws.regions.Regions;
import com.amazonaws.services.quicksight.AmazonQuickSight;
import com.amazonaws.services.quicksight.AmazonQuickSightClientBuilder;
```

```

import
    com.amazonaws.services.quicksight.model.GenerateEmbedUrlForRegisteredUserRequest;
import com.amazonaws.services.quicksight.model.GenerateEmbedUrlForRegisteredUserResult;
import
    com.amazonaws.services.quicksight.model.RegisteredUserEmbeddingExperienceConfiguration;
import
    com.amazonaws.services.quicksight.model.RegisteredUserQuickSightConsoleEmbeddingConfiguration;

/**
 * Class to call QuickSight AWS SDK to get url for QuickSight console embedding.
 */
public class GetQuicksightEmbedUrlRegisteredUserQSConsoleEmbedding {

    private final AmazonQuickSight quickSightClient;

    public GetQuicksightEmbedUrlRegisteredUserQSConsoleEmbedding() {
        this.quickSightClient = AmazonQuickSightClientBuilder
            .standard()
            .withRegion(Regions.US_EAST_1.getName())
            .withCredentials(new AWSCredentialsProvider() {
                @Override
                public AWSCredentials getCredentials() {
                    // provide actual IAM access key and secret key here
                    return new BasicAWSCredentials("access-key", "secret-key");
                }

                @Override
                public void refresh() {
                }
            })
            .build();
    }

    public String getQuicksightEmbedUrl(
        final String accountId,
        final String userArn, // Registered user arn to use for embedding. Refer
        // to Get Embed Url section in developer portal to find out how to get user arn for a
        // QuickSight user.
        final List<String> allowedDomains, // Runtime allowed domain for embedding
        final String initialPath
    ) throws Exception {
        final RegisteredUserEmbeddingExperienceConfiguration experienceConfiguration =
            new RegisteredUserEmbeddingExperienceConfiguration()

```

```

        .withQuickSightConsole(new
RegisteredUserQuickSightConsoleEmbeddingConfiguration().withInitialPath(initialPath));
        final GenerateEmbedUrlForRegisteredUserRequest
generateEmbedUrlForRegisteredUserRequest = new
GenerateEmbedUrlForRegisteredUserRequest();
        generateEmbedUrlForRegisteredUserRequest.setAwsAccountId(accountId);
        generateEmbedUrlForRegisteredUserRequest.setUserArn(userArn);
        generateEmbedUrlForRegisteredUserRequest.setAllowedDomains(allowedDomains);

generateEmbedUrlForRegisteredUserRequest.setExperienceConfiguration(experienceConfiguration);

        final GenerateEmbedUrlForRegisteredUserResult
generateEmbedUrlForRegisteredUserResult =
quickSightClient.generateEmbedUrlForRegisteredUser(generateEmbedUrlForRegisteredUserRequest);

        return generateEmbedUrlForRegisteredUserResult.getEmbedUrl();
    }
}

```

JavaScript

```

global.fetch = require('node-fetch');
const AWS = require('aws-sdk');

function generateEmbedUrlForRegisteredUser(
    accountId,
    dashboardId,
    openIdToken, // Cognito-based token
    userArn, // registered user arn
    roleArn, // IAM user role to use for embedding
    sessionName, // Session name for the roleArn assume role
    allowedDomains, // Runtime allowed domain for embedding
    getEmbedUrlCallback, // GetEmbedUrl success callback method
    errorCallback // GetEmbedUrl error callback method
) {
    const stsClient = new AWS.STS();
    let stsParams = {
        RoleSessionName: sessionName,
        WebIdentityToken: openIdToken,
        RoleArn: roleArn
    }

    stsClient.assumeRoleWithWebIdentity(stsParams, function(err, data) {

```



```

    if (err) {
        console.log('Error assuming role');
        console.log(err, err.stack);
        errorCallback(err);
    } else {
        const getDashboardParams = {
            "AwsAccountId": accountId,
            "ExperienceConfiguration": {
                "QuickSightConsole": {
                    "InitialPath": '/start'
                }
            },
            "UserArn": userArn,
            "AllowedDomains": allowedDomains,
            "SessionLifetimeInMinutes": 600
        };

        const quicksightGetDashboard = new AWS.QuickSight({
            region: process.env.AWS_REGION,
            credentials: {
                accessKeyId: data.Credentials.AccessKeyId,
                secretAccessKey: data.Credentials.SecretAccessKey,
                sessionToken: data.Credentials.SessionToken,
                expiration: data.Credentials.Expiration
            }
        });

        quicksightGetDashboard.generateEmbedUrlForRegisteredUser(getDashboardParams,
        function(err, data) {
            if (err) {
                console.log(err, err.stack);
                errorCallback(err);
            } else {
                const result = {
                    "statusCode": 200,
                    "headers": {
                        "Access-Control-Allow-Origin": "*", // Use your website
                        "Access-Control-Allow-Headers": "Content-Type"
                    },
                    "body": JSON.stringify(data),
                    "isBase64Encoded": false
                }
            }
        });
    }

```

```

        getEmbedUrlCallback(result);
    }
    });
}
});
}

```

Python3

```

import json
import boto3
from botocore.exceptions import ClientError

# Create QuickSight and STS clients
qs = boto3.client('quicksight', region_name='us-east-1')
sts = boto3.client('sts')

# Function to generate embedded URL
# accountId: AWS account ID
# userArn: arn of registered user
# allowedDomains: Runtime allowed domain for embedding
# roleArn: IAM user role to use for embedding
# sessionName: session name for the roleArn assume role
def generateEmbeddingURL(accountId, userArn, allowedDomains, roleArn, sessionName):
    try:
        assumedRole = sts.assume_role(
            RoleArn = roleArn,
            RoleSessionName = sessionName,
        )
    except ClientError as e:
        return "Error assuming role: " + str(e)
    else:
        assumedRoleSession = boto3.Session(
            aws_access_key_id = assumedRole['Credentials']['AccessKeyId'],
            aws_secret_access_key = assumedRole['Credentials']['SecretAccessKey'],
            aws_session_token = assumedRole['Credentials']['SessionToken'],
        )
        try:
            quickSightClient = assumedRoleSession.client('quicksight', region_name='us-
east-1')

            experienceConfiguration = {
                "QuickSightConsole": {

```

```

        "InitialPath": "/start"
    }
}
response = quickSightClient.generate_embed_url_for_registered_user(
    AwsAccountId = accountId,
    ExperienceConfiguration = experienceConfiguration,
    UserArn = userArn,
    AllowedDomains = allowedDomains,
    SessionLifetimeInMinutes = 600
)

return {
    'statusCode': 200,
    'headers': {"Access-Control-Allow-Origin": "*", "Access-Control-Allow-Headers": "Content-Type"},
    'body': json.dumps(response),
    'isBase64Encoded': bool('false')
}
except ClientError as e:
    return "Error generating embedding url: " + str(e)

```

Node.js

The following example shows the JavaScript (Node.js) that you can use on the app server to generate the URL for the embedded console session. You can use this URL in your website or app to display the console session.

Example

```

const AWS = require('aws-sdk');
const https = require('https');

var quicksightClient = new AWS.Service({
    apiConfig: require('./quicksight-2018-04-01.min.json'),
    region: 'us-east-1',
});

quicksightClient.generateEmbedUrlForRegisteredUser({
    'AwsAccountId': '111122223333',
    'ExperienceConfiguration': {
        'QuickSightConsole': {
            'InitialPath': '/start'
        }
    }
}

```

```
    },
    'UserArn': 'REGISTERED_USER_ARN',
    'AllowedDomains': allowedDomains,
    'SessionLifetimeInMinutes': 100
  }, function(err, data) {
    console.log('Errors: ');
    console.log(err);
    console.log('Response: ');
    console.log(data);
  });
```

Example

```
// The URL returned is over 900 characters. For this example, we've shortened the
// string for
// readability and added ellipsis to indicate that it's incomplete.
{
  Status: 200,
  EmbedUrl: 'https://quicksightdomain/embed/12345/dashboards/67890...',
  RequestId: '7bee030e-f191-45c4-97fe-d9faf0e03713'
}
```

.NET/C#

The following example shows the .NET/C# code that you can use on the app server to generate the URL for the embedded console session. You can use this URL in your website or app to display the console.

Example

```
using System;
using Amazon.QuickSight;
using Amazon.QuickSight.Model;

namespace GenerateDashboardEmbedUrlForRegisteredUser
{
    class Program
    {
        static void Main(string[] args)
        {
            var quicksightClient = new AmazonQuickSightClient(
                AccessKey,
                SecretAccessKey,
```

```

        SessionToken,
        Amazon.RegionEndpoint.USEast1);
    try
    {
        RegisteredUserQuickSightConsoleEmbeddingConfiguration
registeredUserQuickSightConsoleEmbeddingConfiguration
        = new RegisteredUserQuickSightConsoleEmbeddingConfiguration
        {
            InitialPath = "/start"
        };
        RegisteredUserEmbeddingExperienceConfiguration
registeredUserEmbeddingExperienceConfiguration
        = new RegisteredUserEmbeddingExperienceConfiguration
        {
            QuickSightConsole =
registeredUserQuickSightConsoleEmbeddingConfiguration
        };

        Console.WriteLine(
            quicksightClient.GenerateEmbedUrlForRegisteredUserAsync(new
GenerateEmbedUrlForRegisteredUserRequest
            {
                AwsAccountId = "111122223333",
                ExperienceConfiguration =
registeredUserEmbeddingExperienceConfiguration,
                UserArn = "REGISTERED_USER_ARN",
                AllowedDomains = allowedDomains,
                SessionLifetimeInMinutes = 100
            }).Result.EmbedUrl
        );
    } catch (Exception ex) {
        Console.WriteLine(ex.Message);
    }
}
}
}
}

```

AWS CLI

To assume the role, choose one of the following AWS Security Token Service (AWS STS) API operations:

- [AssumeRole](#) – Use this operation when you're using an IAM identity to assume the role.

- [AssumeRoleWithWebIdentity](#) – Use this operation when you're using a web identity provider to authenticate your user.
- [AssumeRoleWithSaml](#) – Use this operation when you're using SAML to authenticate your users.

The following example shows the CLI command to set the IAM role. The role needs to have permissions enabled for `quicksight:GenerateEmbedUrlForRegisteredUser`. If you are taking a just-in-time approach to add users when they first open Amazon Quick Sight, the role also needs permissions enabled for `quicksight:RegisterUser`.

```
aws sts assume-role \  
  --role-arn "arn:aws:iam::111122223333:role/embedding_quicksight_dashboard_role" \  
  --role-session-name john.doe@example.com
```

The `assume-role` operation returns three output parameters: the access key, the secret key, and the session token.

Note

If you get an `ExpiredToken` error when calling the `AssumeRole` operation, this is probably because the previous `SESSION_TOKEN` is still in the environment variables. Clear this by setting the following variables:

- `AWS_ACCESS_KEY_ID`
- `AWS_SECRET_ACCESS_KEY`
- `AWS_SESSION_TOKEN`

The following example shows how to set these three parameters in the CLI. If you're using a Microsoft Windows machine, use `set` instead of `export`.

```
export AWS_ACCESS_KEY_ID      = "access_key_from_assume_role"  
export AWS_SECRET_ACCESS_KEY  = "secret_key_from_assume_role"  
export AWS_SESSION_TOKEN      = "session_token_from_assume_role"
```

Running these commands sets the role session ID of the user visiting your website to `embedding_quicksight_console_session_role/john.doe@example.com`. The role session ID is made up of the role name from `role-arn` and the `role-session-name` value. Using the unique role session ID for each user ensures that appropriate permissions are set for each user. It

also prevents any throttling of user access. Throttling is a security feature that prevents the same user from accessing Amazon Quick Sight from multiple locations.

The role session ID also becomes the user name in Amazon Quick Sight. You can use this pattern to provision your users in Amazon Quick Sight ahead of time, or to provision them the first time they access a console session.

The following example shows the CLI command that you can use to provision a user. For more information about [RegisterUser](#), [DescribeUser](#), and other Amazon Quick Sight API operations, see the [Amazon Quick Sight API Reference](#).

```
aws quicksight register-user \  
  --aws-account-id 111122223333 \  
  --namespace default \  
  --identity-type IAM \  
  --iam-arn "arn:aws:iam::111122223333:role/embedding_quicksight_dashboard_role" \  
  --user-role READER \  
  --user-name jhnd \  
  --session-name "john.doe@example.com" \  
  --email john.doe@example.com \  
  --region us-east-1 \  
  --custom-permissions-name TeamA1
```

If the user is authenticated through Microsoft AD, you don't need to use `RegisterUser` to set them up. Instead, they should be automatically subscribed the first time they access Amazon Quick Sight. For Microsoft AD users, you can use `DescribeUser` to get the user ARN.

The first time a user accesses Amazon Quick Sight, you can also add this user to the appropriate group. The following example shows the CLI command to add a user to a group.

```
aws quicksight create-group-membership \  
  --aws-account-id=111122223333 \  
  --namespace=default \  
  --group-name=financeusers \  
  --member-name="embedding_quicksight_dashboard_role/john.doe@example.com"
```

You now have a user of your app who is also a user of Amazon Quick Sight, and who has access to the Amazon Quick Sight console session.

Finally, to get a signed URL for the console session, call `generate-embed-url-for-registered-user` from the app server. This returns the embeddable console session URL. The

following example shows how to generate the URL for an embedded console session using a server-side call for users authenticated through AWS Managed Microsoft AD or single sign-on (IAM Identity Center).

```
aws quicksight generate-embed-url-for-registered-user \
  --aws-account-id 111122223333 \
  --entry-point the-url-for--the-console-session \
  --session-lifetime-in-minutes 600 \
  --user-arn arn:aws:quicksight:us-east-1:111122223333:user/default/
embedding_quicksight_dashboard_role/embeddingsession
--allowed-domains '["domain1","domain2"]' \
--experience-configuration QuickSightConsole={InitialPath="/start"}
```

For more information about using this operation, see [GenerateEmbedUrlForRegisteredUser](#). You can use this and other API operations in your own code.

Step 3: Embed the console session URL

In the following section, you can find out how you can use the [Amazon Quick Sight Embedding SDK](#) (JavaScript) to embed the console session URL from step 3 in your website or application page. With the SDK, you can do the following:

- Place the console session on an HTML page.
- Pass parameters into the console session.
- Handle error states with messages that are customized to your application.

Call the `GenerateEmbedUrlForRegisteredUser` API operation to generate the URL that you can embed in your app. This URL is valid for 5 minutes, and the resulting session is valid for up to 10 hours. The API operation provides the URL with an `auth_code` that enables a single-sign on session.

The following shows an example response from `generate-embed-url-for-registered-user`.

```
//The URL returned is over 900 characters. For this example, we've shortened the string
for
//readability and added ellipsis to indicate that it's incomplete.
{
  "Status": "200",
  "EmbedUrl": "https://quicksightdomain/embedding/12345/start...",
  "RequestId": "7bee030e-f191-45c4-97fe-d9faf0e03713"
```



```
}
```

Embed this console session in your webpage by using the Amazon Quick Sight [Embedding SDK](#) or by adding this URL into an iframe. If you set a fixed height and width number (in pixels), Amazon Quick Sight uses those and doesn't change your visual as your window resizes. If you set a relative percent height and width, Amazon Quick Sight provides a responsive layout that is modified as your window size changes. By using the Amazon Quick Sight Embedding SDK, you can also control parameters within the console session and receive callbacks in terms of page load completion and errors.

The domain that is going to host embedded dashboards must be on the *allow list*, the list of approved domains for your Quick subscription. This requirement protects your data by keeping unapproved domains from hosting embedded dashboards. For more information about adding domains for an embedded console, see [Allow listing domains at runtime with the Amazon Quick Sight API](#).

The following example shows how to use the generated URL. This code is generated on your app server.

SDK 2.0

```
<!DOCTYPE html>
<html>

  <head>
    <title>Console Embedding Example</title>
    <script src="https://unpkg.com/amazon-quicksight-embedding-sdk@2.0.0/dist/
quicksight-embedding-js-sdk.min.js"></script>
    <script type="text/javascript">
      const embedSession = async() => {
        const {
          createEmbeddingContext,
        } = QuickSightEmbedding;

        const embeddingContext = await createEmbeddingContext({
          onChange: (changeEvent, metadata) => {
            console.log('Context received a change', changeEvent,
metadata);
          },
        });

        const frameOptions = {
```

```

        url: "<YOUR_EMBED_URL>", // replace this value with the url
generated via embedding API
        container: '#experience-container',
        height: "700px",
        width: "1000px",
        onChange: (changeEvent, metadata) => {
            switch (changeEvent.eventName) {
                case 'FRAME_MOUNTED': {
                    console.log("Do something when the experience frame is
mounted.");

                    break;
                }
                case 'FRAME_LOADED': {
                    console.log("Do something when the experience frame is
loaded.");

                    break;
                }
            }
        },
    };

    const contentOptions = {
        onMessage: async (messageEvent, experienceMetadata) => {
            switch (messageEvent.eventName) {
                case 'ERROR_OCCURRED': {
                    console.log("Do something when the embedded experience
fails loading.");

                    break;
                }
            }
        }
    };

    const embeddedConsoleExperience = await
embeddingContext.embedConsole(frameOptions, contentOptions);
};
</script>
</head>

<body onload="embedSession()">
    <div id="experience-container"></div>
</body>

</html>

```

SDK 1.0

```
<!DOCTYPE html>
<html>

  <head>
    <title>QuickSight Console Embedding</title>
    <script src="https://unpkg.com/amazon-quicksight-embedding-sdk@1.0.15/dist/
quicksight-embedding-js-sdk.min.js"></script>
    <script type="text/javascript">
      var session

      function onError(payload) {
        console.log("Do something when the session fails loading");
      }

      function embedSession() {
        var containerDiv = document.getElementById("embeddingContainer");
        var options = {
          // replace this dummy url with the one generated via embedding API
          url: "https://us-east-1.quicksight.aws.amazon.com/sn/dashboards/
dashboardId?isauthcode=true&identityprovider=quicksight&code=authcode", // replace this
dummy url with the one generated via embedding API
          container: containerDiv,
          parameters: {
            country: "United States"
          },
          scrolling: "no",
          height: "700px",
          width: "1000px",
          locale: "en-US",
          footerPaddingEnabled: true,
          defaultEmbeddingVisualType: "TABLE", // this option only applies to
QuickSight console embedding and is not used for dashboard embedding
        };
        session = QuickSightEmbedding.embedSession(options);
        session.on("error", onError);
      }

      function onCountryChange(obj) {
        session.setParameters({country: obj.value});
      }
    </script>
  </head>
```

```
<body onload="embedSession()">
  <span>
    <label for="country">Country</label>
    <select id="country" name="country" onchange="onCountryChange(this)">
      <option value="United States">United States</option>
      <option value="Mexico">Mexico</option>
      <option value="Canada">Canada</option>
    </select>
  </span>
  <div id="embeddingContainer"></div>
</body>

</html>
```

For this example to work, make sure to use the Amazon Quick Sight Embedding SDK to load the embedded console session on your website using JavaScript. To get your copy, do one of the following:

- Download the [Amazon Quick Sight Embedding SDK](#) from GitHub. This repository is maintained by a group of Amazon Quick Sight developers.
- Download the latest embedding SDK version from <https://www.npmjs.com/package/amazon-quicksight-embedding-sdk>.
- If you use npm for JavaScript dependencies, download and install it by running the following command.

```
npm install amazon-quicksight-embedding-sdk
```

Enabling Generative BI features in embedded consoles for registered users

Applies to: Enterprise Edition

Intended audience: Amazon Quick developers

You can enable the following Generative BI features in your embedded console:

- **Executive summaries:** When enabled, registered Author Pro and Reader Pro users can generate executive summaries that provide a summary of all insights that Amazon Quick Sight has generated for the dashboard to easily discover key insights.
- **Authoring:** When enabled, Author Pro users can use Generative BI to build calculated fields and build and refine visuals.
- **Q&A:** When enabled, Author Pro and Reader Pro users can use the AI-powered Q&A to both suggest and answer questions related their data.
- **Data stories:** When enabled, Author Pro and Reader Pro users can provide details to quickly generate a first draft of their data story.

To enable Generative BI features in embedded consoles for registered users

- Follow the steps in [Embedding the full functionality of the Amazon Quick Sight console for registered users](#) to embed a console with the following changes:
 - a. When generating the URL in Step 2, set `Enabled: true` in the `FeatureConfigurations` parameter for each of the features you want to enable in the [GenerateEmbedUrlForRegisteredUser](#) or [GenerateEmbedUrlForRegisteredUserWithIdentity](#) APIs, as shown in the following example. If no configuration is provided, the features are disabled by default.

```
ExperienceConfiguration: {
  QuickSightConsole: {
    InitialPath: "initial_path",
    AmazonQInQuickSight: {
      FeatureConfigurations: {
        COMMENT: Enable executive summaries
        ExecutiveSummary: {
          Enabled: true
        },
        COMMENT: Enable Generative BI authoring
        GenerativeAuthoring: {
          Enabled: true
        },
        COMMENT: Enable Q&A
        DataQnA: {
          Enabled: true
        },
        COMMENT: Enable data stories
```

```
DataStories: {  
    Enabled: true  
}  
  
}  
  
}  
  
}  
  
}
```

- b. When embedding the console URL with the Amazon Quick Sight Embedding SDK in Step 3, set the values in the following example as desired. If no configuration is provided, the features are disabled by default.

Note

There is no SDK option for enabling data stories. If data stories are enabled with the API as shown in the previous step, they will be available to registered users.

```
const contentOptions = {
  toolbarOptions: {
    executiveSummary: true, // Enable executive summaries
    buildVisual: true, // Enable Generative BI authoring
    dataQnA: true // Enable Q&A
  }
};
```

Embedding the Amazon Q in Amazon Quick Sight Generative Q&A experience

Intended audience: Amazon Quick developers

In the following sections, you can find detailed information about how to set up an embedded Generative Q&A experience that uses enhanced NLQ capabilities powered by LLMs. The Generative Q&A experience is the recommended replacement for the embedded Q Search Bar and provides an updated BI experience for users.

Topics

- [Embedding the Amazon Q in Amazon Quick Sight Generative Q&A experience for registered users](#)
- [Embedding the Amazon Q in Quick Generative Q&A experience for anonymous \(unregistered\) users](#)

Embedding the Amazon Q in Amazon Quick Sight Generative Q&A experience for registered users

In the following sections, you can find detailed information about how to set up an embedded Generative Q&A experience for registered users of Amazon Quick Sight.

Topics

- [Step 1: Set up permissions](#)
- [Step 2: Generate the URL with the authentication code attached](#)
- [Step 3: Embed the Generative Q&A experience URL](#)
- [Optional embedded Generative Q&A experience functionalities](#)

Step 1: Set up permissions

In the following section, you can find how to set up permissions for your backend application or web server to embed the Generative Q&A experience. This task requires administrative access to AWS Identity and Access Management (IAM).

Each user who accesses a Generative Q&A experience assumes a role that gives them Amazon Quick Sight access and permissions. To make this possible, create an IAM role in your AWS account. Associate an IAM policy with the role to provide permissions to any user who assumes it. The IAM role needs to provide permissions to retrieve embedding URLs for a specific user pool.

With the help of the wildcard character *, you can grant the permissions to generate a URL for all users in a specific namespace. Or you can grant permissions to generate a URL for a subset of users in specific namespaces. For this, you add `quicksight:GenerateEmbedUrlForRegisteredUser`.

You can create a condition in your IAM policy that limits the domains that developers can list in the `AllowedDomains` parameter of a `GenerateEmbedUrlForRegisteredUser` API operation. The `AllowedDomains` parameter is an optional parameter. It grants developers the option to override

the static domains that are configured in the **Manage Amazon Quick Sight** menu and instead list up to three domains or subdomains that can access a generated URL. This URL is then embedded in a developer's website. Only the domains that are listed in the parameter can access the embedded Generative Q&A experience. Without this condition, developers can list any domain on the internet in the `AllowedDomains` parameter.

To limit the domains that developers can use with this parameter, add an `AllowedEmbeddingDomains` condition to your IAM policy. For more information about the `AllowedDomains` parameter, see [GenerateEmbedUrlForRegisteredUser](#) in the *Amazon Quick Sight API Reference*.

Security best practice for IAM condition operators

Improperly configured IAM condition operators can allow unauthorized access to your embedded Quick resources through URL variations. When using the `quicksight:AllowedEmbeddingDomains` condition key in your IAM policies, use condition operators that either allow specific domains or deny all domains that are not specifically allowed. For more information about IAM condition operators, see [IAM JSON policy elements: Condition operators](#) in the IAM User Guide.

Many different URL variations can point to the same resource. For example, the following URLs all resolve to the same content:

- `https://example.com`
- `https://example.com/`
- `https://Example.com`

If your policy uses operators that do not account for these URL variations, an attacker can bypass your restrictions by providing equivalent URL variations.

You must validate that your IAM policy uses appropriate condition operators to prevent bypass vulnerabilities and ensure that only your intended domains can access your embedded resources.

The following sample policy provides these permissions.

Also, if you're creating first-time users who will be Amazon Quick Sight readers, make sure to add the `quicksight:RegisterUser` permission in the policy.

The following sample policy provides permission to retrieve an embedding URL for first-time users who are to be Amazon Quick Sight readers.

Finally, your application's IAM identity must have a trust policy associated with it to allow access to the role that you just created. This means that when a user accesses your application, your application can assume the role on the user's behalf and provision the user in Amazon Quick Sight.

The following example shows a sample trust policy.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowLambdaFunctionsToAssumeThisRole",
      "Effect": "Allow",
      "Principal": {
        "Service": "lambda.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    },
    {
      "Sid": "AllowEC2InstancesToAssumeThisRole",
      "Effect": "Allow",
      "Principal": {
        "Service": "ec2.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

For more information regarding trust policies for OpenID Connect or Security Assertion Markup Language (SAML) authentication, see the following sections of the *IAM User Guide*:

- [Creating a role for web identity or OpenID Connect federation \(console\)](#)
- [Creating a role for SAML 2.0 federation \(console\)](#)

Step 2: Generate the URL with the authentication code attached

In the following section, you can find how to authenticate your user and get the embeddable Q topic URL on your application server. If you plan to embed the Generative Q&A experience for IAM or Amazon Quick Sight identity types, share the Q topic with the users.

When a user accesses your app, the app assumes the IAM role on the user's behalf. Then the app adds the user to Amazon Quick Sight, if that user doesn't already exist. Next, it passes an identifier as the unique role session ID.

Performing the described steps ensures that each viewer of the Q topic is uniquely provisioned in Amazon Quick Sight. It also enforces per-user settings, such as the row-level security and dynamic defaults for parameters. Tag-based row-level security can be used for anonymous user embedding of the Q bar.

The following examples perform the IAM authentication on the user's behalf. This code runs on your app server.

Java

```
import com.amazonaws.auth.AWSCredentials;
import com.amazonaws.auth.BasicAWSCredentials;
import com.amazonaws.auth.AWSCredentialsProvider;
import com.amazonaws.regions.Regions;
import com.amazonaws.services.quicksight.AmazonQuickSight;
import com.amazonaws.services.quicksight.AmazonQuickSightClientBuilder;
import
    com.amazonaws.services.quicksight.model.GenerateEmbedUrlForRegisteredUserRequest;
import com.amazonaws.services.quicksight.model.GenerateEmbedUrlForRegisteredUserResult;
import
    com.amazonaws.services.quicksight.model.RegisteredUserEmbeddingExperienceConfiguration;
import
    com.amazonaws.services.quicksight.model.RegisteredUserGenerativeQnAEmbeddingConfiguration;

/**
 * Class to call QuickSight AWS SDK to get url for embedding Generative Q&A experience.
 */
public class RegisteredUserGenerativeQnAEmbeddingSample {

    private final AmazonQuickSight quickSightClient;

    public RegisteredUserGenerativeQnAEmbeddingSample() {
        this.quickSightClient = AmazonQuickSightClientBuilder
```

```

        .standard()
        .withRegion(Regions.US_EAST_1.getName())
        .withCredentials(new AWSCredentialsProvider() {
            @Override
            public AWSCredentials getCredentials() {
                // provide actual IAM access key and secret key here
                return new BasicAWSCredentials("access-key", "secret-
key");
            }

            @Override
            public void refresh() {
            }
        })
        .build();
    }

    public String getQuicksightEmbedUrl(
        final String accountId, // AWS Account ID
        final String topicId, // Topic ID to embed
        final List<String> allowedDomains, // Runtime allowed domain for embedding
        final String userArn // Registered user arn to use for embedding. Refer to
Get Embed Url section in developer portal to find how to get user arn for a QuickSight
user.
    ) throws Exception {

        final RegisteredUserEmbeddingExperienceConfiguration experienceConfiguration =
new RegisteredUserEmbeddingExperienceConfiguration()
            .withGenerativeQnA(new
RegisteredUserGenerativeQnAEmbeddingConfiguration().withInitialTopicId(topicId));
        final GenerateEmbedUrlForRegisteredUserRequest
generateEmbedUrlForRegisteredUserRequest = new
GenerateEmbedUrlForRegisteredUserRequest();
        generateEmbedUrlForRegisteredUserRequest.setAwsAccountId(accountId);
        generateEmbedUrlForRegisteredUserRequest.setUserArn(userArn);
        generateEmbedUrlForRegisteredUserRequest.setAllowedDomains(allowedDomains);

        generateEmbedUrlForRegisteredUserRequest.setExperienceConfiguration(experienceConfiguration);

        final GenerateEmbedUrlForRegisteredUserResult
generateEmbedUrlForRegisteredUserResult =
quicksightClient.generateEmbedUrlForRegisteredUser(generateEmbedUrlForRegisteredUserRequest);

```

```
        return generateEmbedUrlForRegisteredUserResult.getEmbedUrl();
    }
}
```

JavaScript

Note

Embed URL generation APIs cannot be called from browsers directly. Refer to the Node.JS example instead.

Python3

```
import json
import boto3
from botocore.exceptions import ClientError

sts = boto3.client('sts')

# Function to generate embedded URL
# accountId: AWS account ID
# topicId: Topic ID to embed
# userArn: arn of registered user
# allowedDomains: Runtime allowed domain for embedding
# roleArn: IAM user role to use for embedding
# sessionName: session name for the roleArn assume role
def getEmbeddingURL(accountId, topicId, userArn, allowedDomains, roleArn, sessionName):
    try:
        assumedRole = sts.assume_role(
            RoleArn = roleArn,
            RoleSessionName = sessionName,
        )
    except ClientError as e:
        return "Error assuming role: " + str(e)
    else:
        assumedRoleSession = boto3.Session(
            aws_access_key_id = assumedRole['Credentials']['AccessKeyId'],
            aws_secret_access_key = assumedRole['Credentials']['SecretAccessKey'],
            aws_session_token = assumedRole['Credentials']['SessionToken'],
        )
        try:
```

```

    quicksightClient = assumedRoleSession.client('quicksight', region_name='us-
west-2')

    response = quicksightClient.generate_embed_url_for_registered_user(
        AwsAccountId=accountId,
        ExperienceConfiguration = {
            'GenerativeQnA': {
                'InitialTopicId': topicId
            }
        },
        UserArn = userArn,
        AllowedDomains = allowedDomains,
        SessionLifetimeInMinutes = 600
    )

    return {
        'statusCode': 200,
        'headers': {"Access-Control-Allow-Origin": "*", "Access-Control-Allow-
Headers": "Content-Type"},
        'body': json.dumps(response),
        'isBase64Encoded': bool('false')}
except ClientError as e:
    return "Error generating embedding url: " + str(e)

```

Node.js

The following example shows the JavaScript (Node.js) that you can use on the app server to generate the URL for the embedded dashboard. You can use this URL in your website or app to display the dashboard.

Example

```

const AWS = require('aws-sdk');
const https = require('https');

var quicksightClient = new AWS.Service({
    region: 'us-east-1'
});

quicksightClient.generateEmbedUrlForRegisteredUser({
    'AwsAccountId': '111122223333',
    'ExperienceConfiguration': {
        'GenerativeQnA': {

```

```

        'InitialTopicId': 'U4zJMVZ2n2stZflc80u3iKySEb3BEV6f'
    }
},
'UserArn': 'REGISTERED_USER_ARN',
'AllowedDomains': allowedDomains,
'SessionLifetimeInMinutes': 100
}, function(err, data) {
    console.log('Errors: ');
    console.log(err);
    console.log('Response: ');
    console.log(data);
});

```

.NET/C#

The following example shows the .NET/C# code that you can use on the app server to generate the URL for the embedded Q search bar. You can use this URL in your website or app to display the Q search bar.

Example

```

using System;
using Amazon.QuickSight;
using Amazon.QuickSight.Model;

namespace GenerateGenerativeQnAEmbedUrlForRegisteredUser
{
    class Program
    {
        static void Main(string[] args)
        {
            var quicksightClient = new AmazonQuickSightClient(
                AccessKey,
                SecretAccessKey,
                SessionToken,
                Amazon.RegionEndpoint.USEast1);

            try
            {
                RegisteredUserGenerativeQnAEmbeddingConfiguration
registeredUserGenerativeQnAEmbeddingConfiguration
                    = new RegisteredUserGenerativeQnAEmbeddingConfiguration
                    {
                        InitialTopicId = "U4zJMVZ2n2stZflc80u3iKySEb3BEV6f"

```

```

        };
        RegisteredUserEmbeddingExperienceConfiguration
registeredUserEmbeddingExperienceConfiguration
        = new RegisteredUserEmbeddingExperienceConfiguration
        {
            GenerativeQnA =
registeredUserGenerativeQnAEmbeddingConfiguration
        };

        Console.WriteLine(
            quicksightClient.GenerateEmbedUrlForRegisteredUserAsync(new
GenerateEmbedUrlForRegisteredUserRequest
            {
                AwsAccountId = "111122223333",
                ExperienceConfiguration =
registeredUserEmbeddingExperienceConfiguration,
                UserArn = "REGISTERED_USER_ARN",
                AllowedDomains = allowedDomains,
                SessionLifetimeInMinutes = 100
            }).Result.EmbedUrl
        );
    } catch (Exception ex) {
        Console.WriteLine(ex.Message);
    }
}
}
}
}

```

AWS CLI

To assume the role, choose one of the following AWS Security Token Service (AWS STS) API operations:

- [AssumeRole](#) – Use this operation when you are using an IAM identity to assume the role.
- [AssumeRoleWithWebIdentity](#) – Use this operation when you are using a web identity provider to authenticate your user.
- [AssumeRoleWithSaml](#) – Use this operation when you are using SAML to authenticate your users.

The following example shows the CLI command to set the IAM role. The role needs to have permissions enabled for `quicksight:GenerateEmbedUrlForRegisteredUser`. If you are

taking a just-in-time approach to add users when they use a topic in the Q search bar, the role also needs permissions enabled for `quicksight:RegisterUser`.

```
aws sts assume-role \  
  --role-arn "arn:aws:iam::111122223333:role/  
embedding_quicksight_q_generative_qna_role" \  
  --role-session-name john.doe@example.com
```

The `assume-role` operation returns three output parameters: the access key, the secret key, and the session token.

Note

If you get an `ExpiredToken` error when calling the `AssumeRole` operation, this is probably because the previous `SESSION_TOKEN` is still in the environment variables. Clear this by setting the following variables:

- `AWS_ACCESS_KEY_ID`
- `AWS_SECRET_ACCESS_KEY`
- `AWS_SESSION_TOKEN`

The following example shows how to set these three parameters in the CLI. For a Microsoft Windows machine, use `set` instead of `export`.

```
export AWS_ACCESS_KEY_ID      = "access_key_from_assume_role"  
export AWS_SECRET_ACCESS_KEY  = "secret_key_from_assume_role"  
export AWS_SESSION_TOKEN      = "session_token_from_assume_role"
```

Running these commands sets the role session ID of the user visiting your website to `embedding_quicksight_q_search_bar_role/john.doe@example.com`. The role session ID is made up of the role name from `role-arn` and the `role-session-name` value. Using the unique role session ID for each user ensures that appropriate permissions are set for each user. It also prevents any throttling of user access. *Throttling* is a security feature that prevents the same user from accessing Amazon Quick Sight from multiple locations.

The role session ID also becomes the user name in Amazon Quick Sight. You can use this pattern to provision your users in Amazon Quick Sight ahead of time, or to provision them the first time that they access the Generative Q&A experience.

The following example shows the CLI command that you can use to provision a user. For more information about [RegisterUser](#), [DescribeUser](#), and other Amazon Quick Sight API operations, see the [Amazon Quick Sight API reference](#).

```
aws quicksight register-user \  
  --aws-account-id 111122223333 \  
  --namespace default \  
  --identity-type IAM \  
  --iam-arn "arn:aws:iam::111122223333:role/  
embedding_quicksight_q_generative_qna_role" \  
  --user-role READER \  
  --user-name jhnd \  
  --session-name "john.doe@example.com" \  
  --email john.doe@example.com \  
  --region us-east-1 \  
  --custom-permissions-name TeamA1
```

If the user is authenticated through Microsoft AD, you don't need to use `RegisterUser` to set them up. Instead, they should be automatically subscribed the first time that they access Amazon Quick Sight. For Microsoft AD users, you can use `DescribeUser` to get the user Amazon Resource Name (ARN).

The first time a user accesses Amazon Quick Sight, you can also add this user to the group that the dashboard is shared with. The following example shows the CLI command to add a user to a group.

```
aws quicksight create-group-membership \  
  --aws-account-id 111122223333 \  
  --namespace default \  
  --group-name financeusers \  
  --member-name "embedding_quicksight_q_generative_qna_role/john.doe@example.com"
```

You now have a user of your app who is also a user of Amazon Quick Sight, and who has access to the dashboard.

Finally, to get a signed URL for the dashboard, call `generate-embed-url-for-registered-user` from the app server. This returns the embeddable dashboard URL. The following example shows how to generate the URL for an embedded dashboard using a server-side call for users authenticated through AWS Managed Microsoft AD or single sign-on (IAM Identity Center).

```
aws quicksight generate-embed-url-for-anonymous-user \  
  --aws-account-id 111122223333 \  
  --region us-east-1
```

```
--namespace default-or-something-else \  
--authorized-resource-arns '["topic-arn-topicId1","topic-arn-topicId2"]' \  
--allowed-domains '["domain1","domain2"]' \  
--experience-configuration 'GenerativeQnA={InitialTopicId="topicId1"}' \  
--session-tags '["Key": tag-key-1, "Value": tag-value-1, {"Key": tag-key-1, "Value": tag-value-1}]' \  
--session-lifetime-in-minutes 15
```

For more information about using this operation, see [GenerateEmbedUrlForRegisteredUser](#). You can use this and other API operations in your own code.

Step 3: Embed the Generative Q&A experience URL

In the following section, you can find how to embed the Generative Q&A experience URL in your website or application page. You do this with the [Amazon Quick Sight embedding SDK](#) (JavaScript). With the SDK, you can do the following:

- Place the Generative Q&A experience on an HTML page.
- Customize the layout and appearance of the embedded experience to fit your application needs.
- Handle error states with messages that are customized to your application.

To generate the URL that you can embed in your app, call the `GenerateEmbedUrlForRegisteredUser` API operation. This URL is valid for 5 minutes, and the resulting session is valid for up to 10 hours. The API operation provides the URL with an `auth_code` value that enables a single-sign on session.

The following shows an example response from `generate-embed-url-for-registered-user`.

```
//The URL returned is over 900 characters. For this example, we've shortened the string  
for  
//readability and added ellipsis to indicate that it's incomplete.  
{  
  "Status": "200",  
  "EmbedUrl": "https://quicksightdomain/embedding/12345/q/search...",  
  "RequestId": "7bee030e-f191-45c4-97fe-d9faf0e03713"  
}
```

Embed the Generative Q&A experience in your webpage by using the [Amazon Quick Sight embedding SDK](#) or by adding this URL into an iframe. If you set a fixed height and width number (in pixels), Amazon Quick Sight uses those and doesn't change your visual as your window resizes. If

you set a relative percent height and width, Amazon Quick Sight provides a responsive layout that is modified as your window size changes.

Make sure that the domain to host the embedded Generative Q&A experience is on the *allow list*, the list of approved domains for your Amazon Quick Sight subscription. This requirement protects your data by keeping unapproved domains from hosting embedded dashboards. For more information about adding domains for an embedded Generative Q&A experience, see [Managing domains](#).

You can use the Amazon Quick Sight Embedding SDK to customize the layout and appearance of the embedded Generative Q&A experience to fit your application. Use the `panelType` property to configure the landing state of the Generative Q&A experience when it renders in your application. Set the `panelType` property to 'FULL' to render the full Generative Q&A experience panel. This panel resembles the experience that Amazon Quick Sight users have in the Amazon Quick Sight console. The frame height of the panel is not changed based on user interaction and respects the value that you set in the `frameOptions.height` property. The image below shows the Generative Q&A experience panel that renders when you set the `panelType` value to 'FULL'.

Set the `panelType` property to 'SEARCH_BAR' to render the Generative Q&A experience as a search bar. This search bar resembles the way that the Q Search Bar renders when it is embedded into an application. The Generative Q&A search bar expands to a larger panel that displays topic selection options, the question suggestion list, the answer panel or the pinboard.

The default minimum height of the Generative Q&A search bar is rendered when the embedded asset loads. It is recommended that you set the `frameOptions.height` value to "38px" to optimize the search bar experience. Use the `focusedHeight` property to set the optimal size of the topic selection dropdown and the question suggestion list. Use the `expandedHeight` property to set the optimal size of the answer panel and pinboard. If you choose the 'SEARCH_BAR' option, it is recommended that you style the parent container with `position: absolute` to avoid unwanted content shifting in your application. The image below shows the Generative Q&A experience search bar that renders when you set the `panelType` value to 'SEARCH_BAR'.

After you configure the `panelType` property, use the Amazon Quick Sight embedding SDK to customize the following properties of the Generative Q&A experience.

- The title of the Generative Q&A panel (Applies only to the `panelType`: FULL option).
- The search bar's placeholder text.
- Whether topic selection is allowed.

- Whether topic names are shown or hidden.
- Whether the Amazon Q icon is shown or hidden (Applies only to the `panelType`: FULL option).
- Whether the pinboard is shown or hidden.
- Whether users can maximize the Generative Q&A panel to fullscreen.
- The theme of the Generative Q&A panel. A custom theme ARN can be passed in the SDK to change the appearance of the frame's content. Amazon Quick Sight starter themes are not supported for embedded Generative BI panels. To use an Amazon Quick Sight starter theme, save it as a custom theme in Amazon Quick Sight.

When you use the Amazon Quick Sight Embedding SDK, the Generative Q&A experience on your page is dynamically resized based on the state. By using the Amazon Quick Sight Embedding SDK, you can also control parameters within the Generative Q&A experience and receive callbacks in terms of page load completion, state changes, and errors.

The following example shows how to use the generated URL. This code is generated on your app server.

SDK 2.0

```
<!DOCTYPE html>
<html>
  <head>
    <title>Generative Q&A Embedding Example</title>
    <script src="https://unpkg.com/amazon-quicksight-embedding-sdk@2.7.0/dist/
quicksight-embedding-js-sdk.min.js"></script>
    <script type="text/javascript">
      const embedGenerativeQnA = async() => {
        const {createEmbeddingContext} = QuickSightEmbedding;

        const embeddingContext = await createEmbeddingContext({
          onChange: (changeEvent, metadata) => {
            console.log('Context received a change', changeEvent,
metadata);
          },
        });

        const frameOptions = {
          url: "<YOUR_EMBED_URL>", // replace this value with the url
generated via embedding API
          container: '#experience-container',
```

```

        height: "700px",
        width: "1000px",
        onChange: (changeEvent, metadata) => {
            switch (changeEvent.eventName) {
                case 'FRAME_MOUNTED': {
                    console.log("Do something when the experience frame is
mounted.");

                    break;
                }
                case 'FRAME_LOADED': {
                    console.log("Do something when the experience frame is
loaded.");

                    break;
                }
            }
        },
    };

    const contentOptions = {
        // Optional panel settings. Default behavior is equivalent to
        {panelType: 'FULL'}
        panelOptions: {
            panelType: 'FULL',
            title: 'custom title', // Optional
            showQIcon: false, // Optional, Default: true
        },
        // Use SEARCH_BAR panel type for the landing state to be similar to
embedQSearchBar
        // with generative capability enabled topics
        /*
        panelOptions: {
            panelType: 'SEARCH_BAR',
            focusedHeight: '250px',
            expandedHeight: '500px',
        },
        */
        showTopicName: false, // Optional, Default: true
        showPinboard: false, // Optional, Default: true
        allowTopicSelection: false, // Optional, Default: true
        allowFullscreen: false, // Optional, Default: true
        searchPlaceholderText: "custom search placeholder", // Optional
        themeOptions: { // Optional
            themeArn: 'arn:aws:quicksight:<Region>:<AWS-Account-ID>:theme/
<Theme-ID>'

```

```

    }
    onMessage: async (messageEvent, experienceMetadata) => {
      switch (messageEvent.eventName) {
        case 'Q_SEARCH_OPENED': {
          // called when pinboard is shown / visuals are rendered
          console.log("Do something when SEARCH_BAR type panel is
expanded");

          break;
        }
        case 'Q_SEARCH_FOCUSED': {
          // called when question suggestions or topic selection
dropdown are shown

          console.log("Do something when SEARCH_BAR type panel is
focused");

          break;
        }
        case 'Q_SEARCH_CLOSED': {
          // called when shrunk to initial bar height
          console.log("Do something when SEARCH_BAR type panel is
collapsed");

          break;
        }
        case 'Q_PANEL_ENTERED_FULLSCREEN': {
          console.log("Do something when panel enters full screen
mode");

          break;
        }
        case 'Q_PANEL_EXITED_FULLSCREEN': {
          console.log("Do something when panel exits full screen
mode");

          break;
        }
        case 'CONTENT_LOADED': {
          console.log("Do something after experience is loaded");
          break;
        }
        case 'ERROR_OCCURRED': {
          console.log("Do something when experience fails to
load");

          break;
        }
      }
    }
  };

```

```
        const embeddedGenerativeQnExperience = await
embeddingContext.embedGenerativeQnA(frameOptions, contentOptions);
    };
</script>
</head>

<body onload="embedGenerativeQnA()">
    <div id="experience-container"></div>
</body>

</html>
```

For this example to work, make sure to use the Amazon Quick Sight Embedding SDK to load the embedded Generative Q&A experience on your website with JavaScript. To get your copy, do one of the following:

- Download the [Amazon Quick Sight embedding SDK](#) from GitHub. This repository is maintained by a group of Amazon Quick Sight developers.
- Download the latest embedding SDK version from <https://www.npmjs.com/package/amazon-quicksight-embedding-sdk>.
- If you use npm for JavaScript dependencies, download and install it by running the following command.

```
npm install amazon-quicksight-embedding-sdk
```

Optional embedded Generative Q&A experience functionalities

The following optional functionalities are available for the embedded Generative Q&A experience with the embedding SDK.

Invoke Generative Q&A search bar actions

- Set a question — This feature sends a question to the Generative Q&A experience and immediately queries the question.

```
embeddedGenerativeQnExperience.setQuestion('show me monthly revenue');
```

- Close the answer panel (applies to the Generative Q&A search bar option) — This feature closes the answer panel and returns the iframe to the original search bar state.

```
embeddedGenerativeQnExperience.close();
```

For more information, see the [Amazon Quick Sight embedding SDK](#).

Embedding the Amazon Q in Quick Generative Q&A experience for anonymous (unregistered) users

Intended audience: Amazon Quick developers

In the following sections, you can find detailed information about how to set up an embedded Generative Q&A experience for anonymous (unregistered) users.

Topics

- [Step 1: Set up permissions](#)
- [Step 2: Generate the URL with the authentication code attached](#)
- [Step 3: Embed the Generative Q&A experience URL](#)
- [Optional embedded Generative Q&A experience functionalities](#)

Step 1: Set up permissions

In the following section, you can find how to set up permissions for your backend application or web server to embed the Generative Q&A experience. This task requires administrative access to AWS Identity and Access Management (IAM).

Each user who accesses a Generative Q&A experience assumes a role that gives them Amazon Quick Sight access and permissions. To make this possible, create an IAM role in your AWS account. Associate an IAM policy with the role to provide permissions to any user who assumes it. The IAM role needs to provide permissions to retrieve embedding URLs for a specific user pool.

With the help of the wildcard character *, you can grant the permissions to generate a URL for all users in a specific namespace. Or you can grant permissions to generate a URL for a subset of users in specific namespaces. For this, you add `quicksight:GenerateEmbedUrlForAnonymousUser`.

You can create a condition in your IAM policy that limits the domains that developers can list in the `AllowedDomains` parameter of a `GenerateEmbedUrlForAnonymousUser` API operation.

The `AllowedDomains` parameter is an optional parameter. It grants developers the option to override the static domains that are configured in the **Manage Amazon Quick Sight** menu and instead list up to three domains or subdomains that can access a generated URL. This URL is then embedded in a developer's website. Only the domains that are listed in the parameter can access the embedded Q search bar. Without this condition, developers can list any domain on the internet in the `AllowedDomains` parameter.

To limit the domains that developers can use with this parameter, add an `AllowedEmbeddingDomains` condition to your IAM policy. For more information about the `AllowedDomains` parameter, see [GenerateEmbedUrlForAnonymousUser](#) in the *Amazon Quick Sight API Reference*.

Security best practice for IAM condition operators

Improperly configured IAM condition operators can allow unauthorized access to your embedded Quick resources through URL variations. When using the `quicksight:AllowedEmbeddingDomains` condition key in your IAM policies, use condition operators that either allow specific domains or deny all domains that are not specifically allowed. For more information about IAM condition operators, see [IAM JSON policy elements: Condition operators](#) in the IAM User Guide.

Many different URL variations can point to the same resource. For example, the following URLs all resolve to the same content:

- `https://example.com`
- `https://example.com/`
- `https://Example.com`

If your policy uses operators that do not account for these URL variations, an attacker can bypass your restrictions by providing equivalent URL variations.

You must validate that your IAM policy uses appropriate condition operators to prevent bypass vulnerabilities and ensure that only your intended domains can access your embedded resources.

Your application's IAM identity must have a trust policy associated with it to allow access to the role that you just created. This means that when a user accesses your application, your application

can assume the role on the user's behalf to load the Generative Q&A experience. The following example shows a sample trust policy.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowLambdaFunctionsToAssumeThisRole",
      "Effect": "Allow",
      "Principal": {
        "Service": "lambda.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    },
    {
      "Sid": "AllowEC2InstancesToAssumeThisRole",
      "Effect": "Allow",
      "Principal": {
        "Service": "ec2.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

For more information regarding trust policies, see [Temporary security credentials in IAM](#) in the *IAM User Guide*

Step 2: Generate the URL with the authentication code attached

In the following section, you can find how to authenticate your user and get the embeddable Q topic URL on your application server.

When a user accesses your app, the app assumes the IAM role on the user's behalf. Then the app adds the user to Amazon Quick Sight, if that user doesn't already exist. Next, it passes an identifier as the unique role session ID.

Java

```
import java.util.List;
import com.amazonaws.auth.AWSCredentials;
import com.amazonaws.auth.AWSCredentialsProvider;
import com.amazonaws.auth.BasicAWSCredentials;
import com.amazonaws.regions.Regions;
import com.amazonaws.services.quicksight.AmazonQuickSight;
import com.amazonaws.services.quicksight.AmazonQuickSightClientBuilder;
import
    com.amazonaws.services.quicksight.model.AnonymousUserGenerativeQnAEmbeddingConfiguration;
import
    com.amazonaws.services.quicksight.model.AnonymousUserEmbeddingExperienceConfiguration;
import com.amazonaws.services.quicksight.model.GenerateEmbedUrlForAnonymousUserRequest;
import com.amazonaws.services.quicksight.model.GenerateEmbedUrlForAnonymousUserResult;
import com.amazonaws.services.quicksight.model.SessionTag;

/**
 * Class to call QuickSight AWS SDK to generate embed url for anonymous user.
 */
public class GenerateEmbedUrlForAnonymousUserExample {

    private final AmazonQuickSight quickSightClient;

    public GenerateEmbedUrlForAnonymousUserExample() {
        quickSightClient = AmazonQuickSightClientBuilder
            .standard()
            .withRegion(Regions.US_EAST_1.getName())
            .withCredentials(new AWSCredentialsProvider() {
                @Override
                public AWSCredentials getCredentials() {
                    // provide actual IAM access key and secret key here
                    return new BasicAWSCredentials("access-key", "secret-key");
                }

                @Override
                public void refresh() {
                }
            })
            .build();
    }

    public String GenerateEmbedUrlForAnonymousUser(
```

```

        final String accountId, // YOUR AWS ACCOUNT ID
        final String initialTopicId, // Q TOPIC ID TO WHICH THE CONSTRUCTED URL POINTS
        AND EXPERIENCE PREPOPULATES INITIALLY
        final String namespace, // ANONYMOUS EMBEDDING REQUIRES SPECIFYING A VALID
        NAMESPACE FOR WHICH YOU WANT THE EMBEDDING URL
        final List<String> authorizedResourceArns, // Q TOPIC ARN LIST TO EMBED
        final List<String> allowedDomains, // RUNTIME ALLOWED DOMAINS FOR EMBEDDING
        final List<SessionTag> sessionTags // SESSION TAGS USED FOR ROW-LEVEL SECURITY
    ) throws Exception {
        AnonymousUserEmbeddingExperienceConfiguration experienceConfiguration = new
        AnonymousUserEmbeddingExperienceConfiguration();
        AnonymousUserGenerativeQnAEmbeddingConfiguration generativeQnAConfiguration =
        new AnonymousUserGenerativeQnAEmbeddingConfiguration();
        generativeQnAConfiguration.setInitialTopicId(initialTopicId);
        experienceConfiguration.setGenerativeQnA(generativeQnAConfiguration);

        GenerateEmbedUrlForAnonymousUserRequest generateEmbedUrlForAnonymousUserRequest
        = new GenerateEmbedUrlForAnonymousUserRequest()
            .withAwsAccountId(accountId)
            .withNamespace(namespace)
            .withAuthorizedResourceArns(authorizedResourceArns)
            .withExperienceConfiguration(experienceConfiguration)
            .withSessionTags(sessionTags)
            .withSessionLifetimeInMinutes(600L); // OPTIONAL: VALUE CAN BE [15-600].
        DEFAULT: 600
            .withAllowedDomains(allowedDomains);

        GenerateEmbedUrlForAnonymousUserResult result =
        quickSightClient.generateEmbedUrlForAnonymousUser(generateEmbedUrlForAnonymousUserRequest);

        return result.getEmbedUrl();
    }
}

```

JavaScript

Note

Embed URL generation APIs cannot be called from browsers directly. Refer to the Node.JS example instead.

Python3

```
import json
import boto3
from botocore.exceptions import ClientError
import time

# Create QuickSight and STS clients
quicksightClient = boto3.client('quicksight', region_name='us-west-2')
sts = boto3.client('sts')

# Function to generate embedded URL for anonymous user
# accountId: YOUR AWS ACCOUNT ID
# topicId: Topic ID to embed
# quicksightNamespace: VALID NAMESPACE WHERE YOU WANT TO DO NOAUTH EMBEDDING
# authorizedResourceArns: TOPIC ARN LIST TO EMBED
# allowedDomains: RUNTIME ALLOWED DOMAINS FOR EMBEDDING
# sessionTags: SESSION TAGS USED FOR ROW-LEVEL SECURITY
def generateEmbedUrlForAnonymousUser(accountId, quicksightNamespace,
    authorizedResourceArns, allowedDomains, sessionTags):
    try:
        response = quicksightClient.generate_embed_url_for_anonymous_user(
            AwsAccountId = accountId,
            Namespace = quicksightNamespace,
            AuthorizedResourceArns = authorizedResourceArns,
            AllowedDomains = allowedDomains,
            ExperienceConfiguration = {
                'GenerativeQnA': {
                    'InitialTopicId': topicId
                }
            },
            SessionTags = sessionTags,
            SessionLifetimeInMinutes = 600
        )

        return {
            'statusCode': 200,
            'headers': {"Access-Control-Allow-Origin": "*", "Access-Control-Allow-Headers": "Content-Type"},
            'body': json.dumps(response),
            'isBase64Encoded': bool('false')
        }
    except ClientError as e:
        print(e)
```

```
return "Error generating embeddedURL: " + str(e)
```

Node.js

The following example shows the JavaScript (Node.js) that you can use on the app server to generate the URL for the embedded dashboard. You can use this URL in your website or app to display the dashboard.

Example

```
const AWS = require('aws-sdk');
const https = require('https');

var quicksightClient = new AWS.Service({
    region: 'us-east-1',
});

quicksightClient.generateEmbedUrlForAnonymousUser({
    'AwsAccountId': '111122223333',
    'Namespace': 'DEFAULT',
    'AuthorizedResourceArns': '["topic-arn-topicId1","topic-arn-topicId2"]',
    'AllowedDomains': allowedDomains,
    'ExperienceConfiguration': {
        'GenerativeQnA': {
            'InitialTopicId': 'U4zJMVZ2n2stZf1c80u3iKySEb3BEV6f'
        }
    },
    'SessionTags': '["Key": tag-key-1,"Value": tag-value-1,{"Key": tag-key-1,"Value": tag-value-1}]',
    'SessionLifetimeInMinutes': 15
}, function(err, data) {
    console.log('Errors: ');
    console.log(err);
    console.log('Response: ');
    console.log(data);
});
```

.NET/C#

The following example shows the .NET/C# code that you can use on the app server to generate the URL for the embedded Q search bar. You can use this URL in your website or app to display the Q search bar.

Example

```
using System;
using Amazon.QuickSight;
using Amazon.QuickSight.Model;

namespace GenerateGenerativeQnAEmbedUrlForAnonymousUser
{
    class Program
    {
        static void Main(string[] args)
        {
            var quicksightClient = new AmazonQuickSightClient(
                AccessKey,
                SecretAccessKey,
                SessionToken,
                Amazon.RegionEndpoint.USEast1);
            try
            {
                AnonymousUserGenerativeQnAEmbeddingConfiguration
anonymousUserGenerativeQnAEmbeddingConfiguration
                = new AnonymousUserGenerativeQnAEmbeddingConfiguration
                {
                    InitialTopicId = "U4zJMVZ2n2stZflc80u3iKySEb3BEV6f"
                };
                AnonymousUserEmbeddingExperienceConfiguration
anonymousUserEmbeddingExperienceConfiguration
                = new AnonymousUserEmbeddingExperienceConfiguration
                {
                    GenerativeQnA =
anonymousUserGenerativeQnAEmbeddingConfiguration
                };

                Console.WriteLine(
                    quicksightClient.GenerateEmbedUrlForAnonymousUserAsync(new
GenerateEmbedUrlForAnonymousUserRequest
                {
                    AwsAccountId = "111122223333",
                    Namespace = "DEFAULT",
                    AuthorizedResourceArns ["topic-arn-topicId1","topic-arn-
topicId2"]',
                    AllowedDomains = allowedDomains,
                    ExperienceConfiguration =
anonymousUserEmbeddingExperienceConfiguration,
```

```

        SessionTags = '[{"Key": tag-key-1,"Value": tag-value-1,{"Key":
tag-key-1,"Value": tag-value-1}}]',
        SessionLifetimeInMinutes = 15,
    })).Result.EmbedUrl
    );
} catch (Exception ex) {
    Console.WriteLine(ex.Message);
}
}
}
}
}

```

AWS CLI

To assume the role, choose one of the following AWS Security Token Service (AWS STS) API operations:

- [AssumeRole](#) – Use this operation when you are using an IAM identity to assume the role.
- [AssumeRoleWithWebIdentity](#) – Use this operation when you are using a web identity provider to authenticate your user.
- [AssumeRoleWithSaml](#) – Use this operation when you are using SAML to authenticate your users.

The following example shows the CLI command to set the IAM role. The role needs to have permissions enabled for `quicksight:GenerateEmbedUrlForAnonymousUser`.

```

aws sts assume-role \
    --role-arn "arn:aws:iam::111122223333:role/
embedding_quicksight_generative_qna_role" \
    --role-session-name anonymous caller

```

The `assume-role` operation returns three output parameters: the access key, the secret key, and the session token.

Note

If you get an `ExpiredToken` error when calling the `AssumeRole` operation, this is probably because the previous `SESSION_TOKEN` is still in the environment variables. Clear this by setting the following variables:

- `AWS_ACCESS_KEY_ID`

- `AWS_SECRET_ACCESS_KEY`
- `AWS_SESSION_TOKEN`

The following example shows how to set these three parameters in the CLI. For a Microsoft Windows machine, use `set` instead of `export`.

```
export AWS_ACCESS_KEY_ID      = "access_key_from_assume_role"
export AWS_SECRET_ACCESS_KEY  = "secret_key_from_assume_role"
export AWS_SESSION_TOKEN      = "session_token_from_assume_role"
```

Running these commands sets the role session ID of the user visiting your website to `embedding_quicksight_q_search_bar_role/QuickSightEmbeddingAnonymousPolicy`. The role session ID is made up of the role name from `role-arn` and the `role-session-name` value. Using the unique role session ID for each user ensures that appropriate permissions are set for each user. It also prevents any throttling of user access. *Throttling* is a security feature that prevents the same user from accessing Amazon Quick Sight from multiple locations. In addition, it keeps each session separate and distinct. If you're using an array of web servers, for example for load balancing, and a session is reconnected to a different server, a new session begins.

To get a signed URL for the dashboard, call `generate-embed-url-for-anonymous-user` from the app server. This returns the embeddable dashboard URL. The following example shows how to generate the URL for an embedded dashboard using a server-side call for users who are making anonymous visits to your web portal or app.

```
aws quicksight generate-embed-url-for-anonymous-user \
--aws-account-id 111122223333 \
--namespace default-or-something-else \
--authorized-resource-arns '["topic-arn-topicId","topic-arn-topicId2"]' \
--allowed-domains '["domain1","domain2"]' \
--experience-configuration 'GenerativeQnA={InitialTopicId="topicId1"}' \
--session-tags '["Key": tag-key-1,"Value": tag-value-1,{"Key": tag-key-1,"Value": tag-value-1}]' \
--session-lifetime-in-minutes 15
```

For more information about using this operation, see [GenerateEmbedUrlForAnonymousUser](#). You can use this and other API operations in your own code.

Step 3: Embed the Generative Q&A experience URL

In the following section, you can find how to embed the Generative Q&A experience URL in your website or application page. You do this with the [Amazon Quick Sight embedding SDK](#) (JavaScript). With the SDK, you can do the following:

- Place the Generative Q&A experience on an HTML page.
- Customize the layout and appearance of the embedded experience to fit your application needs.
- Handle error states with messages that are customized to your application.

To generate the URL that you can embed in your app, call the `GenerateEmbedUrlForAnonymousUser` API operation. This URL is valid for 5 minutes, and the resulting session is valid for up to 10 hours. The API operation provides the URL with an `auth_code` value that enables a single-sign on session.

The following shows an example response from `generate-embed-url-for-anonymous-user`.

```
//The URL returned is over 900 characters. For this example, we've shortened the string
for
//readability and added ellipsis to indicate that it's incomplete.{
  "Status": "200",
  "EmbedUrl": "https://quicksightdomain/embedding/12345/q/search...",
  "RequestId": "7bee030e-f191-45c4-97fe-d9faf0e03713"
}
```

Embed the Generative Q&A experience in your webpage with the [Amazon Quick Sight embedding SDK](#) or by adding this URL into an `iframe`. If you set a fixed height and width number (in pixels), Amazon Quick Sight uses those and doesn't change your visual as your window resizes. If you set a relative percent height and width, Amazon Quick Sight provides a responsive layout that is modified as your window size changes.

Make sure that the domain to host the Generative Q&A experience is on the *allow list*, the list of approved domains for your Amazon Quick Sight subscription. This requirement protects your data by keeping unapproved domains from hosting embedded Generative Q&A experiences. For more information about adding domains for an embedded Generative Q&A experience, see [Managing domains](#).

You can use the Amazon Quick Sight Embedding SDK to customize the layout and appearance of the embedded Generative Q&A experience to fit your application. Use the `panelType` property to

configure the landing state of the Generative Q&A experience when it renders in your application. Set the `panelType` property to `'FULL'` to render the full Generative Q&A experience panel. This panel resembles the experience that Amazon Quick Sight users have in the Amazon Quick Sight console. The frame height of the panel is not changed based on user interaction and respects the value that you set in the `frameOptions.height` property. The image below shows the Generative Q&A experience panel that renders when you set the `panelType` value to `'FULL'`.

Set the `panelType` property to `'SEARCH_BAR'` to render the Generative Q&A experience as a search bar. This search bar resembles the way that the Q Search Bar renders when it is embedded into an application. The Generative Q&A search bar expands to a larger panel that displays topic selection options, the question suggestion list, the answer panel or the pinboard.

The default minimum height of the Generative Q&A search bar is rendered when the embedded asset loads. It is recommended that you set the `frameOptions.height` value to `"38px"` to optimize the search bar experience. Use the `focusedHeight` property to set the optimal size of the topic selection dropdown and the question suggestion list. Use the `expandedHeight` property to set the optimal size of the answer panel and pinboard. If you choose the `'SEARCH_BAR'` option, it is recommended that you style the parent container with `position; absolute` to avoid unwanted content shifting in your application. The image below shows the Generative Q&A experience search bar that renders when you set the `panelType` value to `'SEARCH_BAR'`.

After you configure the `panelType` property, use the Amazon Quick Sight embedding SDK to customize the following properties of the Generative Q&A experience.

- The title of the Generative Q&A panel (Applies only to the `panelType: FULL` option).
- The search bar's placeholder text.
- Whether topic selection is allowed.
- Whether topic names are shown or hidden.
- Whether the Amazon Q icon is shown or hidden (Applies only to the `panelType: FULL` option).
- Whether the pinboard is shown or hidden.
- Whether users can maximize the Generative Q&A panel to fullscreen.
- The theme of the Generative Q&A panel. A custom theme ARN can be passed in the SDK to change the appearance of the frame's content. Amazon Quick Sight starter themes are not supported for embedded Generative BI panels. To use a Amazon Quick Sight starter theme, save it as a custom theme in Amazon Quick Sight.

When you use the Amazon Quick Sight Embedding SDK, the Generative Q&A experience on your page is dynamically resized based on the state. With the Amazon Quick Sight Embedding SDK, you can also control parameters within the Generative Q&A experience and receive callbacks in terms of page load completion, state changes, and errors.

The following example shows how to use the generated URL. This code is generated on your app server.

SDK 2.0

```
<!DOCTYPE html>
<html>
  <head>
    <title>Generative Q&A Embedding Example</title>
    <script src="https://unpkg.com/amazon-quicksight-embedding-sdk@2.7.0/dist/
quicksight-embedding-js-sdk.min.js"></script>
    <script type="text/javascript">
      const embedGenerativeQnA = async() => {
        const {createEmbeddingContext} = QuickSightEmbedding;

        const embeddingContext = await createEmbeddingContext({
          onChange: (changeEvent, metadata) => {
            console.log('Context received a change', changeEvent,
metadata);
          },
        });

        const frameOptions = {
          url: "<YOUR_EMBED_URL>", // replace this value with the url
generated via embedding API
          container: '#experience-container',
          height: "700px",
          width: "1000px",
          onChange: (changeEvent, metadata) => {
            switch (changeEvent.eventName) {
              case 'FRAME_MOUNTED': {
                console.log("Do something when the experience frame is
mounted.");
                break;
              }
              case 'FRAME_LOADED': {
                console.log("Do something when the experience frame is
loaded.");
            }
```

```

        break;
    }
    },
};

const contentOptions = {
    // Optional panel settings. Default behavior is equivalent to
    {panelType: 'FULL'}
    panelOptions: {
        panelType: 'FULL',
        title: 'custom title', // Optional
        showQIcon: false, // Optional, Default: true
    },
    // Use SEARCH_BAR panel type for the landing state to be similar to
    embedQSearchBar
    // with generative capability enabled topics
    /*
    panelOptions: {
        panelType: 'SEARCH_BAR',
        focusedHeight: '250px',
        expandedHeight: '500px',
    },
    */
    showTopicName: false, // Optional, Default: true
    showPinboard: false, // Optional, Default: true
    allowTopicSelection: false, // Optional, Default: true
    allowFullscreen: false, // Optional, Default: true
    searchPlaceholderText: "custom search placeholder", // Optional
    themeOptions: { // Optional
        themeArn: 'arn:aws:quicksight:<Region>:<AWS-Account-ID>:theme/
<Theme-ID>'
    }
    onMessage: async (messageEvent, experienceMetadata) => {
        switch (messageEvent.eventName) {
            case 'Q_SEARCH_OPENED': {
                // called when pinboard is shown / visuals are rendered
                console.log("Do something when SEARCH_BAR type panel is
expanded");

                break;
            }
            case 'Q_SEARCH_FOCUSED': {
                // called when question suggestions or topic selection
dropdown are shown

```

```

        console.log("Do something when SEARCH_BAR type panel is
focused");
        break;
    }
    case 'Q_SEARCH_CLOSED': {
        // called when shrunk to initial bar height
        console.log("Do something when SEARCH_BAR type panel is
collapsed");
        break;
    }
    case 'Q_PANEL_ENTERED_FULLSCREEN': {
        console.log("Do something when panel enters full screen
mode");
        break;
    }
    case 'Q_PANEL_EXITED_FULLSCREEN': {
        console.log("Do something when panel exits full screen
mode");
        break;
    }
    case 'CONTENT_LOADED': {
        console.log("Do something after experience is loaded");
        break;
    }
    case 'ERROR_OCCURRED': {
        console.log("Do something when experience fails to
load");
        break;
    }
}
}
};
const embeddedGenerativeQnExperience = await
embeddingContext.embedGenerativeQnA(frameOptions, contentOptions);
};
</script>
</head>

<body onload="embedGenerativeQnA()">
    <div id="experience-container"></div>
</body>

</html>

```

For this example to work, make sure to use the Amazon Quick Sight Embedding SDK to load the embedded Generative Q&A experience on your website with JavaScript. To get your copy, do one of the following:

- Download the [Amazon Quick Sight embedding SDK](#) from GitHub. This repository is maintained by a group of Amazon Quick Sight developers.
- Download the latest embedding SDK version from <https://www.npmjs.com/package/amazon-quicksight-embedding-sdk>.
- If you use npm for JavaScript dependencies, download and install it by running the following command.

```
npm install amazon-quicksight-embedding-sdk
```

Optional embedded Generative Q&A experience functionalities

The following optional functionalities are available for the embedded Generative Q&A experience with the embedding SDK.

Invoke Generative Q&A search bar actions

- Set a question — This feature sends a question to the Generative Q&A experience and immediately queries the question.

```
embeddedGenerativeQnExperience.setQuestion('show me monthly revenue');
```

- Close the answer panel (applies to the Generative Q&A search bar option) — This feature closes the answer panel and returns the iframe to the original search bar state.

```
embeddedGenerativeQnExperience.close();
```

For more information, see the [Amazon Quick Sight embedding SDK](#).

Embedding the Amazon Quick Sight Q search bar (Classic)

Intended audience: Amazon Quick developers

Note

The embedded Amazon Quick Sight Q search bar provides the classic Amazon Quick Sight Q&A experience. Amazon Quick Sight integrates with Amazon Q Business to launch a new Generative Q&A experience. Developers are recommended to use the new Generative Q&A experience. For more information on the embedded Generative Q&A experience, see [Embedding the Amazon Q in Amazon Quick Sight Generative Q&A experience](#).

Use the following topics to learn about embedding the Amazon Quick Sight Q search bar with the Amazon Quick Sight APIs.

Topics

- [Embedding the Amazon Quick Sight Q search bar for registered users](#)
- [Embedding the Amazon Quick Sight Q search bar for anonymous \(unregistered\) users](#)

Embedding the Amazon Quick Sight Q search bar for registered users

Applies to: Enterprise Edition

Intended audience: Amazon Quick developers

Note

The embedded Amazon Quick Sight Q search bar provides the classic Amazon Quick Sight Q&A experience. Amazon Quick Sight integrates with Amazon Q Business to launch a new Generative Q&A experience. Developers are recommended to use the new Generative Q&A experience. For more information on the embedded Generative Q&A experience, see [Embedding the Amazon Q in Amazon Quick Sight Generative Q&A experience](#).

In the following sections, you can find detailed information about how to set up an embedded Amazon Quick Sight Q search bar for registered users of Amazon Quick Sight.

Topics

- [Step 1: Set up permissions](#)
- [Step 2: Generate the URL with the authentication code attached](#)
- [Step 3: Embed the Q search bar URL](#)
- [Optional Amazon Quick Sight Q search bar embedding functionalities](#)

Step 1: Set up permissions

Note

The embedded Amazon Quick Sight Q search bar provides the classic Amazon Quick Sight Q&A experience. Amazon Quick Sight integrates with Amazon Q Business to launch a new Generative Q&A experience. Developers are recommended to use the new Generative Q&A experience. For more information on the embedded Generative Q&A experience, see [Embedding the Amazon Q in Amazon Quick Sight Generative Q&A experience](#).

In the following section, you can find how to set up permissions for your backend application or web server to embed the Q search bar. This task requires administrative access to AWS Identity and Access Management (IAM).

Each user who accesses a dashboard assumes a role that gives them Amazon Quick Sight access and permissions to the dashboard. To make this possible, create an IAM role in your AWS account. Associate an IAM policy with the role to provide permissions to any user who assumes it. The IAM role needs to provide permissions to retrieve embedding URLs for a specific user pool.

With the help of the wildcard character *, you can grant the permissions to generate a URL for all users in a specific namespace. Or you can grant permissions to generate a URL for a subset of users in specific namespaces. For this, you add `quicksight:GenerateEmbedUrlForRegisteredUser`.

You can create a condition in your IAM policy that limits the domains that developers can list in the `AllowedDomains` parameter of a `GenerateEmbedUrlForRegisteredUser` API operation. The `AllowedDomains` parameter is an optional parameter. It grants developers the option to override the static domains that are configured in the **Manage Amazon Quick Sight** menu and instead list up to three domains or subdomains that can access a generated URL. This URL is then embedded in a developer's website. Only the domains that are listed in the parameter can access

the embedded Q search bar. Without this condition, developers can list any domain on the internet in the `AllowedDomains` parameter.

To limit the domains that developers can use with this parameter, add an `AllowedEmbeddingDomains` condition to your IAM policy. For more information about the `AllowedDomains` parameter, see [GenerateEmbedUrlForRegisteredUser](#) in the *Amazon Quick Sight API Reference*.

Security best practice for IAM condition operators

Improperly configured IAM condition operators can allow unauthorized access to your embedded Quick resources through URL variations. When using the `quicksight:AllowedEmbeddingDomains` condition key in your IAM policies, use condition operators that either allow specific domains or deny all domains that are not specifically allowed. For more information about IAM condition operators, see [IAM JSON policy elements: Condition operators](#) in the IAM User Guide.

Many different URL variations can point to the same resource. For example, the following URLs all resolve to the same content:

- `https://example.com`
- `https://example.com/`
- `https://Example.com`

If your policy uses operators that do not account for these URL variations, an attacker can bypass your restrictions by providing equivalent URL variations.

You must validate that your IAM policy uses appropriate condition operators to prevent bypass vulnerabilities and ensure that only your intended domains can access your embedded resources.

The following sample policy provides these permissions.

Also, if you're creating first-time users who will be Amazon Quick Sight readers, make sure to add the `quicksight:RegisterUser` permission in the policy.

The following sample policy provides permission to retrieve an embedding URL for first-time users who are to be Amazon Quick Sight readers.

Finally, your application's IAM identity must have a trust policy associated with it to allow access to the role that you just created. This means that when a user accesses your application, your application can assume the role on the user's behalf and provision the user in Amazon Quick Sight.

The following example shows a sample trust policy.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowLambdaFunctionsToAssumeThisRole",
      "Effect": "Allow",
      "Principal": {
        "Service": "lambda.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    },
    {
      "Sid": "AllowEC2InstancesToAssumeThisRole",
      "Effect": "Allow",
      "Principal": {
        "Service": "ec2.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

For more information regarding trust policies for OpenID Connect or Security Assertion Markup Language (SAML) authentication, see the following sections of the *IAM User Guide*:

- [Creating a role for web identity or OpenID Connect federation \(console\)](#)
- [Creating a role for SAML 2.0 federation \(console\)](#)

Step 2: Generate the URL with the authentication code attached

Note

The embedded Amazon Quick Sight Q search bar provides the classic Amazon Quick Sight Q&A experience. Amazon Quick Sight integrates with Amazon Q Business to launch a new Generative Q&A experience. Developers are recommended to use the new Generative Q&A experience. For more information on the embedded Generative Q&A experience, see [Embedding the Amazon Q in Amazon Quick Sight Generative Q&A experience](#).

In the following section, you can find how to authenticate your user and get the embeddable Q topic URL on your application server. If you plan to embed the Q bar for IAM or Amazon Quick Sight identity types, share the Q topic with the users.

When a user accesses your app, the app assumes the IAM role on the user's behalf. Then the app adds the user to Amazon Quick Sight, if that user doesn't already exist. Next, it passes an identifier as the unique role session ID.

Performing the described steps ensures that each viewer of the Q topic is uniquely provisioned in Amazon Quick Sight. It also enforces per-user settings, such as the row-level security and dynamic defaults for parameters.

The following examples perform the IAM authentication on the user's behalf. This code runs on your app server.

Java

```
import com.amazonaws.auth.AWSCredentials;
import com.amazonaws.auth.BasicAWSCredentials;
import com.amazonaws.auth.AWSCredentialsProvider;
import com.amazonaws.regions.Regions;
import com.amazonaws.services.quicksight.AmazonQuickSight;
import com.amazonaws.services.quicksight.AmazonQuickSightClientBuilder;
import
    com.amazonaws.services.quicksight.model.GenerateEmbedUrlForRegisteredUserRequest;
import com.amazonaws.services.quicksight.model.GenerateEmbedUrlForRegisteredUserResult;
import
    com.amazonaws.services.quicksight.model.RegisteredUserEmbeddingExperienceConfiguration;
import
    com.amazonaws.services.quicksight.model.RegisteredUserQSearchBarEmbeddingConfiguration;
```

```

    /**
     * Class to call QuickSight AWS SDK to get url for embedding the Q search bar.
     */
    public class RegisteredUserQSearchBarEmbeddingConfiguration {

        private final AmazonQuickSight quickSightClient;

        public RegisteredUserQSearchBarEmbeddingConfiguration() {
            this.quickSightClient = AmazonQuickSightClientBuilder
                .standard()
                .withRegion(Regions.US_EAST_1.getName())
                .withCredentials(new AWSCredentialsProvider() {
                    @Override
                    public AWSCredentials getCredentials() {
                        // provide actual IAM access key and secret key here
                        return new BasicAWSCredentials("access-key", "secret-
key");
                    }

                    @Override
                    public void refresh() {
                    }
                })
                .build();
        }

        public String getQuicksightEmbedUrl(
            final String accountId, // AWS Account ID
            final String topicId, // Topic ID to embed
            final List<String> allowedDomains, // Runtime allowed domain for embedding
            final String userArn // Registered user arn to use for embedding. Refer to
            Get Embed Url section in developer portal to find how to get user arn for a QuickSight
            user.
        ) throws Exception {
            final RegisteredUserEmbeddingExperienceConfiguration experienceConfiguration =
            new RegisteredUserEmbeddingExperienceConfiguration()
                .withQSearchBar(new
                RegisteredUserQSearchBarEmbeddingConfiguration().withInitialTopicId(topicId));
            final GenerateEmbedUrlForRegisteredUserRequest
            generateEmbedUrlForRegisteredUserRequest = new
            GenerateEmbedUrlForRegisteredUserRequest();
            generateEmbedUrlForRegisteredUserRequest.setAwsAccountId(accountId);

```

```

        generateEmbedUrlForRegisteredUserRequest.setUserArn(userArn);
        generateEmbedUrlForRegisteredUserRequest.setAllowedDomains(allowedDomains);

generateEmbedUrlForRegisteredUserRequest.setExperienceConfiguration(QSearchBar);

        final GenerateEmbedUrlForRegisteredUserResult
generateEmbedUrlForRegisteredUserResult =
quickSightClient.generateEmbedUrlForRegisteredUser(generateEmbedUrlForRegisteredUserRequest);

        return generateEmbedUrlForRegisteredUserResult.getEmbedUrl();
    }
}

```

JavaScript

```

global.fetch = require('node-fetch');
const AWS = require('aws-sdk');

function generateEmbedUrlForRegisteredUser(
    accountId,
    topicId, // Topic ID to embed
    openIdToken, // Cognito-based token
    userArn, // registered user arn
    roleArn, // IAM user role to use for embedding
    sessionName, // Session name for the roleArn assume role
    allowedDomains, // Runtime allowed domain for embedding
    getEmbedUrlCallback, // GetEmbedUrl success callback method
    errorCallback // GetEmbedUrl error callback method
) {
    const stsClient = new AWS.STS();
    let stsParams = {
        RoleSessionName: sessionName,
        WebIdentityToken: openIdToken,
        RoleArn: roleArn
    }

    stsClient.assumeRoleWithWebIdentity(stsParams, function(err, data) {
        if (err) {
            console.log('Error assuming role');
            console.log(err, err.stack);
            errorCallback(err);
        } else {
            const getQSearchBarParams = {

```

```

    "AwsAccountId": accountId,
    "ExperienceConfiguration": {
      "QSearchBar": {
        "InitialTopicId": topicId
      }
    },
    "UserArn": userArn,
    "AllowedDomains": allowedDomains,
    "SessionLifetimeInMinutes": 600
  };

  const quicksightGetQSearchBar = new AWS.QuickSight({
    region: process.env.AWS_REGION,
    credentials: {
      accessKeyId: data.Credentials.AccessKeyId,
      secretAccessKey: data.Credentials.SecretAccessKey,
      sessionToken: data.Credentials.SessionToken,
      expiration: data.Credentials.Expiration
    }
  });

  quicksightGetQSearchBar.generateEmbedUrlForRegisteredUser(getQSearchBarParams,
  function(err, data) {
    if (err) {
      console.log(err, err.stack);
      errorCallback(err);
    } else {
      const result = {
        "statusCode": 200,
        "headers": {
          "Access-Control-Allow-Origin": "*", // Use your website
          "Access-Control-Allow-Headers": "Content-Type"
        },
        "body": JSON.stringify(data),
        "isBase64Encoded": false
      }
      getEmbedUrlCallback(result);
    }
  });
});

```

```
}

```

Python3

```
import json
import boto3
from botocore.exceptions import ClientError

sts = boto3.client('sts')

# Function to generate embedded URL
# accountId: AWS account ID
# topicId: Topic ID to embed
# userArn: arn of registered user
# allowedDomains: Runtime allowed domain for embedding
# roleArn: IAM user role to use for embedding
# sessionName: session name for the roleArn assume role
def getEmbeddingURL(accountId, topicId, userArn, allowedDomains, roleArn, sessionName):
    try:
        assumedRole = sts.assume_role(
            RoleArn = roleArn,
            RoleSessionName = sessionName,
        )
    except ClientError as e:
        return "Error assuming role: " + str(e)
    else:
        assumedRoleSession = boto3.Session(
            aws_access_key_id = assumedRole['Credentials']['AccessKeyId'],
            aws_secret_access_key = assumedRole['Credentials']['SecretAccessKey'],
            aws_session_token = assumedRole['Credentials']['SessionToken'],
        )
        try:
            quicksightClient = assumedRoleSession.client('quicksight', region_name='us-
west-2')
            response = quicksightClient.generate_embed_url_for_registered_user(
                AwsAccountId=accountId,
                ExperienceConfiguration = {
                    "QSearchBar": {
                        "InitialTopicId": topicId
                    }
                },
                UserArn = userArn,
                AllowedDomains = allowedDomains,
```



```
        SessionLifetimeInMinutes = 600
    )

    return {
        'statusCode': 200,
        'headers': {"Access-Control-Allow-Origin": "*", "Access-Control-Allow-Headers": "Content-Type"},
        'body': json.dumps(response),
        'isBase64Encoded': bool('false')
    }
except ClientError as e:
    return "Error generating embedding url: " + str(e)
```

Node.js

The following example shows the JavaScript (Node.js) that you can use on the app server to generate the URL for the embedded dashboard. You can use this URL in your website or app to display the dashboard.

Example

```
const AWS = require('aws-sdk');
const https = require('https');

var quicksightClient = new AWS.Service({
    apiConfig: require('./quicksight-2018-04-01.min.json'),
    region: 'us-east-1',
});

quicksightClient.generateEmbedUrlForRegisteredUser({
    'AwsAccountId': '111122223333',
    'ExperienceConfiguration': {
        'QSearchBar': {
            'InitialTopicId': 'U4zJMVZ2n2stZflc80u3iKySEb3BEV6f'
        }
    },
    'UserArn': 'REGISTERED_USER_ARN',
    'AllowedDomains': allowedDomains,
    'SessionLifetimeInMinutes': 100
}, function(err, data) {
    console.log('Errors: ');
    console.log(err);
    console.log('Response: ');
```

```
console.log(data);
});
```

Example

```
//The URL returned is over 900 characters. For this example, we've shortened the string
//for
//readability and added ellipsis to indicate that it's incomplete.
{
    Status: 200,
    EmbedUrl: "https://quicksightdomain/embed/12345/dashboards/67890/sheets/12345/
visuals/67890...",
    RequestId: '7bee030e-f191-45c4-97fe-d9faf0e03713'
}
```

.NET/C#

The following example shows the .NET/C# code that you can use on the app server to generate the URL for the embedded Q search bar. You can use this URL in your website or app to display the Q search bar.

Example

```
using System;
using Amazon.QuickSight;
using Amazon.QuickSight.Model;

namespace GenerateDashboardEmbedUrlForRegisteredUser
{
    class Program
    {
        static void Main(string[] args)
        {
            var quicksightClient = new AmazonQuickSightClient(
                AccessKey,
                SecretAccessKey,
                SessionToken,
                Amazon.RegionEndpoint.USEast1);
            try
            {
                RegisteredUserQSearchBarEmbeddingConfiguration
registeredUserQSearchBarEmbeddingConfiguration
                    = new RegisteredUserQSearchBarEmbeddingConfiguration
```

```

        {
            InitialTopicId = "U4zJMVZ2n2stZflc80u3iKySEb3BEV6f"
        };
        RegisteredUserEmbeddingExperienceConfiguration
registeredUserEmbeddingExperienceConfiguration
        = new RegisteredUserEmbeddingExperienceConfiguration
        {
            QSearchBar = registeredUserQSearchBarEmbeddingConfiguration
        };

        Console.WriteLine(
            quicksightClient.GenerateEmbedUrlForRegisteredUserAsync(new
GenerateEmbedUrlForRegisteredUserRequest
            {
                AwsAccountId = "111122223333",
                ExperienceConfiguration =
registeredUserEmbeddingExperienceConfiguration,
                UserArn = "REGISTERED_USER_ARN",
                AllowedDomains = allowedDomains,
                SessionLifetimeInMinutes = 100
            }).Result.EmbedUrl
        );
    } catch (Exception ex) {
        Console.WriteLine(ex.Message);
    }
}
}
}
}
}

```

AWS CLI

To assume the role, choose one of the following AWS Security Token Service (AWS STS) API operations:

- [AssumeRole](#) – Use this operation when you are using an IAM identity to assume the role.
- [AssumeRoleWithWebIdentity](#) – Use this operation when you are using a web identity provider to authenticate your user.
- [AssumeRoleWithSaml](#) – Use this operation when you are using SAML to authenticate your users.

The following example shows the CLI command to set the IAM role. The role needs to have permissions enabled for `quicksight:GenerateEmbedUrlForRegisteredUser`. If you are

taking a just-in-time approach to add users when they use a topic in the Q search bar, the role also needs permissions enabled for `quicksight:RegisterUser`.

```
aws sts assume-role \  
  --role-arn "arn:aws:iam::111122223333:role/embedding_quicksight_q_search_bar_role" \  
  --role-session-name john.doe@example.com
```

The `assume-role` operation returns three output parameters: the access key, the secret key, and the session token.

Note

If you get an `ExpiredToken` error when calling the `AssumeRole` operation, this is probably because the previous `SESSION_TOKEN` is still in the environment variables. Clear this by setting the following variables:

- `AWS_ACCESS_KEY_ID`
- `AWS_SECRET_ACCESS_KEY`
- `AWS_SESSION_TOKEN`

The following example shows how to set these three parameters in the CLI. For a Microsoft Windows machine, use `set` instead of `export`.

```
export AWS_ACCESS_KEY_ID      = "access_key_from_assume_role" \  
export AWS_SECRET_ACCESS_KEY  = "secret_key_from_assume_role" \  
export AWS_SESSION_TOKEN      = "session_token_from_assume_role"
```

Running these commands sets the role session ID of the user visiting your website to `embedding_quicksight_q_search_bar_role/john.doe@example.com`. The role session ID is made up of the role name from `role-arn` and the `role-session-name` value. Using the unique role session ID for each user ensures that appropriate permissions are set for each user. It also prevents any throttling of user access. *Throttling* is a security feature that prevents the same user from accessing Amazon Quick Sight from multiple locations.

The role session ID also becomes the user name in Amazon Quick Sight. You can use this pattern to provision your users in Amazon Quick Sight ahead of time, or to provision them the first time that they access the Q search bar.

The following example shows the CLI command that you can use to provision a user. For more information about [RegisterUser](#), [DescribeUser](#), and other Amazon Quick Sight API operations, see the [Amazon Quick Sight API reference](#).

```
aws quicksight register-user \  
  --aws-account-id 111122223333 \  
  --namespace default \  
  --identity-type IAM \  
  --iam-arn "arn:aws:iam::111122223333:role/embedding_quicksight_q_search_bar_role" \  
  --user-role READER \  
  --user-name jhnd \  
  --session-name "john.doe@example.com" \  
  --email john.doe@example.com \  
  --region us-east-1 \  
  --custom-permissions-name TeamA1
```

If the user is authenticated through Microsoft AD, you don't need to use `RegisterUser` to set them up. Instead, they should be automatically subscribed the first time that they access Amazon Quick Sight. For Microsoft AD users, you can use `DescribeUser` to get the user Amazon Resource Name (ARN).

The first time a user accesses Amazon Quick Sight, you can also add this user to the group that the dashboard is shared with. The following example shows the CLI command to add a user to a group.

```
aws quicksight create-group-membership \  
  --aws-account-id=111122223333 \  
  --namespace=default \  
  --group-name=financeusers \  
  --member-name="embedding_quicksight_q_search_bar_role/john.doe@example.com"
```

You now have a user of your app who is also a user of Amazon Quick Sight, and who has access to the dashboard.

Finally, to get a signed URL for the dashboard, call `generate-embed-url-for-registered-user` from the app server. This returns the embeddable dashboard URL. The following example shows how to generate the URL for an embedded dashboard using a server-side call for users authenticated through AWS Managed Microsoft AD or single sign-on (IAM Identity Center).

```
aws quicksight generate-embed-url-for-registered-user \  
  --aws-account-id 111122223333 \  
  --session-lifetime-in-minutes 600 \  
  --region us-east-1
```

```
--user-arn arn:aws:quicksight:us-east-1:111122223333:user/default/
embedding_quicksight_q_search_bar_role/embeddingsession
--allowed-domains '["domain1", "domain2"]' \
--experience-configuration QSearchBar={InitialTopicId=U4zJMVZ2n2stZflc80u3iKySEb3BEV6f}
```

For more information about using this operation, see [GenerateEmbedUrlForRegisteredUser](#). You can use this and other API operations in your own code.

Step 3: Embed the Q search bar URL

Note

The embedded Amazon Quick Sight Q search bar provides the classic Amazon Quick Sight Q&A experience. Amazon Quick Sight integrates with Amazon Q Business to launch a new Generative Q&A experience. Developers are recommended to use the new Generative Q&A experience. For more information on the embedded Generative Q&A experience, see [Embedding the Amazon Q in Amazon Quick Sight Generative Q&A experience](#).

In the following section, you can find how to embed the Q search bar URL from step 3 in your website or application page. You do this with the [Amazon Quick Sight embedding SDK](#) (JavaScript). With the SDK, you can do the following:

- Place the Q search bar on an HTML page.
- Pass parameters into the Q search bar.
- Handle error states with messages that are customized to your application.

To generate the URL that you can embed in your app, call the `GenerateEmbedUrlForRegisteredUser` API operation. This URL is valid for 5 minutes, and the resulting session is valid for up to 10 hours. The API operation provides the URL with an `auth_code` value that enables a single-sign on session.

The following shows an example response from `generate-embed-url-for-registered-user`.

```
//The URL returned is over 900 characters. For this example, we've shortened the string
for
//readability and added ellipsis to indicate that it's incomplete.
{
    "Status": "200",
```

```
"EmbedUrl": "https://quicksightdomain/embedding/12345/q/search...",  
"RequestId": "7bee030e-f191-45c4-97fe-d9faf0e03713"  
}
```

Embed the Q search bar in your webpage by using the [Amazon Quick Sight embedding SDK](#) or by adding this URL into an iframe. If you set a fixed height and width number (in pixels), Amazon Quick Sight uses those and doesn't change your visual as your window resizes. If you set a relative percent height and width, Amazon Quick Sight provides a responsive layout that is modified as your window size changes.

To do this, make sure that the domain to host the embedded Q search bar is on the *allow list*, the list of approved domains for your Amazon Quick Sight subscription. This requirement protects your data by keeping unapproved domains from hosting embedded dashboards. For more information about adding domains for an embedded Q search bar, see [Managing domains and embedding](#).

When you use the Amazon Quick Sight Embedding SDK, the Q search bar on your page is dynamically resized based on the state. By using the Amazon Quick Sight Embedding SDK, you can also control parameters within the Q search bar and receive callbacks in terms of page load completion and errors.

The following example shows how to use the generated URL. This code is generated on your app server.

SDK 2.0

```
<!DOCTYPE html>  
<html>  
  
  <head>  
    <title>Q Search Bar Embedding Example</title>  
    <script src="https://unpkg.com/amazon-quicksight-embedding-sdk@2.0.0/dist/  
quicksight-embedding-js-sdk.min.js"></script>  
    <script type="text/javascript">  
      const embedQSearchBar = async() => {  
        const {  
          createEmbeddingContext,  
        } = QuickSightEmbedding;  
  
        const embeddingContext = await createEmbeddingContext({  
          onChange: (changeEvent, metadata) => {  
            console.log('Context received a change', changeEvent,  
metadata);
```

```

        },
    });

    const frameOptions = {
        url: "<YOUR_EMBED_URL>", // replace this value with the url
generated via embedding API
        container: '#experience-container',
        height: "700px",
        width: "1000px",
        onChange: (changeEvent, metadata) => {
            switch (changeEvent.eventName) {
                case 'FRAME_MOUNTED': {
                    console.log("Do something when the experience frame is
mounted.");

                    break;
                }
                case 'FRAME_LOADED': {
                    console.log("Do something when the experience frame is
loaded.");

                    break;
                }
            }
        },
    };

    const contentOptions = {
        hideTopicName: false,
        theme: '<YOUR_THEME_ID>',
        allowTopicSelection: true,
        onMessage: async (messageEvent, experienceMetadata) => {
            switch (messageEvent.eventName) {
                case 'Q_SEARCH_OPENED': {
                    console.log("Do something when Q Search content
expanded");

                    break;
                }
                case 'Q_SEARCH_CLOSED': {
                    console.log("Do something when Q Search content
collapsed");

                    break;
                }
                case 'Q_SEARCH_SIZE_CHANGED': {
                    console.log("Do something when Q Search size changed");
                    break;
                }
            }
        }
    };

```



```

        }
        case 'CONTENT_LOADED': {
            console.log("Do something when the Q Search is
loaded.");

            break;
        }
        case 'ERROR_OCCURRED': {
            console.log("Do something when the Q Search fails
loading.");

            break;
        }
    }
}
};
const embeddedDashboardExperience = await
embeddingContext.embedQSearchBar(frameOptions, contentOptions);
};
</script>
</head>

<body onload="embedQSearchBar()">
    <div id="experience-container"></div>
</body>

</html>

```

SDK 1.0

```

<!DOCTYPE html>
<html>

    <head>
        <title>QuickSight Q Search Bar Embedding</title>
        <script src="https://unpkg.com/amazon-quicksight-embedding-sdk@1.18.0/dist/
quicksight-embedding-js-sdk.min.js"></script>
        <script type="text/javascript">
            var session

            function onError(payload) {
                console.log("Do something when the session fails loading");
            }

            function onOpen() {

```

```

        console.log("Do something when the Q search bar opens");
    }

    function onClose() {
        console.log("Do something when the Q search bar closes");
    }

    function embedQSearchBar() {
        var containerDiv = document.getElementById("embeddingContainer");
        var options = {
            url: "https://us-east-1.quicksight.aws.amazon.com/sn/dashboards/
dashboardId?isauthcode=true&identityprovider=quicksight&code=authcode", // replace this
dummy url with the one generated via embedding API
            container: containerDiv,
            width: "1000px",
            locale: "en-US",
            qSearchBarOptions: {
                expandCallback: onOpen,
                collapseCallback: onClose,
                iconDisabled: false,
                topicNameDisabled: false,
                themeId: 'bdb844d0-0fe9-4d9d-b520-0fe602d93639',
                allowTopicSelection: true
            }
        };
        session = QuickSightEmbedding.embedQSearchBar(options);
        session.on("error", onError);
    }

    function onCountryChange(obj) {
        session.setParameters({country: obj.value});
    }
</script>
</head>

<body onload="embedQSearchBar()">
    <div id="embeddingContainer"></div>
</body>

</html>

```

For this example to work, make sure to use the Amazon Quick Sight Embedding SDK to load the embedded dashboard on your website using JavaScript. To get your copy, do one of the following:

- Download the [Amazon Quick Sight embedding SDK](#) from GitHub. This repository is maintained by a group of Amazon Quick Sight developers.
- Download the latest embedding SDK version from <https://www.npmjs.com/package/amazon-quicksight-embedding-sdk>.
- If you use npm for JavaScript dependencies, download and install it by running the following command.

```
npm install amazon-quicksight-embedding-sdk
```

Optional Amazon Quick Sight Q search bar embedding functionalities

Note

The embedded Amazon Quick Sight Q search bar provides the classic Amazon Quick Sight Q&A experience. Amazon Quick Sight integrates with Amazon Q Business to launch a new Generative Q&A experience. Developers are recommended to use the new Generative Q&A experience. For more information on the embedded Generative Q&A experience, see [Embedding the Amazon Q in Amazon Quick Sight Generative Q&A experience](#).

The following optional functionalities are available for the embedded Q search bar using the embedding SDK.

Invoke Q search bar actions

The following options are only supported for Q search bar embedding.

- Set a Q search bar question — This feature sends a question to the Q search bar and immediately queries the question. It also automatically opens the Q popover.

```
qBar.setQBarQuestion('show me monthly revenue');
```

- Close the Q popover — This feature closes the Q popover and returns the iframe to the original Q search bar size.

```
qBar.closeQPopover();
```

For more information, see the [Amazon Quick Sight embedding SDK](#).

Embedding the Amazon Quick Sight Q search bar for anonymous (unregistered) users

Intended audience: Amazon Quick developers

Note

The embedded Amazon Quick Sight Q search bar provides the classic Amazon Quick Sight Q&A experience. Amazon Quick Sight integrates with Amazon Q Business to launch a new Generative Q&A experience. Developers are recommended to use the new Generative Q&A experience. For more information on the embedded Generative Q&A experience, see [Embedding the Amazon Q in Amazon Quick Sight Generative Q&A experience](#).

In the following sections, you can find detailed information about how to set up an embedded Amazon Quick Sight Q search bar for anonymous (unregistered) users.

Topics

- [Step 1: Set up permissions](#)
- [Step 2: Generate the URL with the authentication code attached](#)
- [Step 3: Embed the Q search bar URL](#)
- [Optional Amazon Quick Sight Q search bar embedding functionalities](#)

Step 1: Set up permissions

Note

The embedded Amazon Quick Sight Q search bar provides the classic Amazon Quick Sight Q&A experience. Amazon Quick Sight integrates with Amazon Q Business to launch a new Generative Q&A experience. Developers are recommended to use the new Generative Q&A experience. For more information on the embedded Generative Q&A experience, see [Embedding the Amazon Q in Amazon Quick Sight Generative Q&A experience](#).

In the following section, you can find how to set up permissions for your backend application or web server to embed the Q search bar. This task requires administrative access to AWS Identity and Access Management (IAM).

Each user who accesses a Q search bar assumes a role that gives them Amazon Quick Sight access and permissions to the Q search bar. To make this possible, create an IAM role in your AWS account. Associate an IAM policy with the role to provide permissions to any user who assumes it. The IAM role needs to provide permissions to retrieve embedding URLs for a specific user pool.

With the help of the wildcard character *, you can grant the permissions to generate a URL for all users in a specific namespace. Or you can grant permissions to generate a URL for a subset of users in specific namespaces. For this, you add `quicksight:GenerateEmbedUrlForAnonymousUser`.

You can create a condition in your IAM policy that limits the domains that developers can list in the `AllowedDomains` parameter of a `GenerateEmbedUrlForAnonymousUser` API operation. The `AllowedDomains` parameter is an optional parameter. It grants developers the option to override the static domains that are configured in the **Manage Amazon Quick Sight** menu and instead list up to three domains or subdomains that can access a generated URL. This URL is then embedded in a developer's website. Only the domains that are listed in the parameter can access the embedded Q search bar. Without this condition, developers can list any domain on the internet in the `AllowedDomains` parameter.

To limit the domains that developers can use with this parameter, add an `AllowedEmbeddingDomains` condition to your IAM policy. For more information about the `AllowedDomains` parameter, see [GenerateEmbedUrlForAnonymousUser](#) in the *Amazon Quick Sight API Reference*.

Security best practice for IAM condition operators

Improperly configured IAM condition operators can allow unauthorized access to your embedded Quick resources through URL variations. When using the `quicksight:AllowedEmbeddingDomains` condition key in your IAM policies, use condition operators that either allow specific domains or deny all domains that are not specifically allowed. For more information about IAM condition operators, see [IAM JSON policy elements: Condition operators](#) in the IAM User Guide.

Many different URL variations can point to the same resource. For example, the following URLs all resolve to the same content:

- `https://example.com`

- <https://example.com/>
- <https://Example.com>

If your policy uses operators that do not account for these URL variations, an attacker can bypass your restrictions by providing equivalent URL variations.

You must validate that your IAM policy uses appropriate condition operators to prevent bypass vulnerabilities and ensure that only your intended domains can access your embedded resources.

Your application's IAM identity must have a trust policy associated with it to allow access to the role that you just created. This means that when a user accesses your application, your application can assume the role on the user's behalf to open the Q search bar. The following example shows a sample trust policy.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowLambdaFunctionsToAssumeThisRole",
      "Effect": "Allow",
      "Principal": {
        "Service": "lambda.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    },
    {
      "Sid": "AllowEC2InstancesToAssumeThisRole",
      "Effect": "Allow",
      "Principal": {
        "Service": "ec2.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

For more information regarding trust policies, see [Temporary security credentials in IAM](#) in the *IAM User Guide*

Step 2: Generate the URL with the authentication code attached

Note

The embedded Amazon Quick Sight Q search bar provides the classic Amazon Quick Sight Q&A experience. Amazon Quick Sight integrates with Amazon Q Business to launch a new Generative Q&A experience. Developers are recommended to use the new Generative Q&A experience. For more information on the embedded Generative Q&A experience, see [Embedding the Amazon Q in Amazon Quick Sight Generative Q&A experience](#).

In the following section, you can find how to authenticate your user and get the embeddable Q topic URL on your application server.

When a user accesses your app, the app assumes the IAM role on the user's behalf. Then the app adds the user to Amazon Quick Sight, if that user doesn't already exist. Next, it passes an identifier as the unique role session ID.

For more information, see [AnonymousUserQSearchBarEmbeddingConfiguration](#).

Java

```
import java.util.List;
import com.amazonaws.auth.AWSCredentials;
import com.amazonaws.auth.AWSCredentialsProvider;
import com.amazonaws.auth.BasicAWSCredentials;
import com.amazonaws.regions.Regions;
import com.amazonaws.services.quicksight.AmazonQuickSight;
import com.amazonaws.services.quicksight.AmazonQuickSightClientBuilder;
import
com.amazonaws.services.quicksight.model.AnonymousUserQSearchBarEmbeddingConfiguration;
import
com.amazonaws.services.quicksight.model.AnonymousUserEmbeddingExperienceConfiguration;
import
com.amazonaws.services.quicksight.model.GenerateEmbedUrlForAnonymousUserRequest;
import
com.amazonaws.services.quicksight.model.GenerateEmbedUrlForAnonymousUserResult;
import com.amazonaws.services.quicksight.model.SessionTag;
```

```

/**
 * Class to call QuickSight AWS SDK to generate embed url for anonymous user.
 */
public class GenerateEmbedUrlForAnonymousUserExample {

    private final AmazonQuickSight quickSightClient;

    public GenerateEmbedUrlForAnonymousUserExample() {
        quickSightClient = AmazonQuickSightClientBuilder
            .standard()
            .withRegion(Regions.US_EAST_1.getName())
            .withCredentials(new AWSCredentialsProvider() {
                @Override
                public AWSCredentials getCredentials() {
                    // provide actual IAM access key and secret key here
                    return new BasicAWSCredentials("access-key", "secret-
key");
                }

                @Override
                public void refresh() {
                }
            })
            .build();
    }

    public String GenerateEmbedUrlForAnonymousUser(
        final String accountId, // YOUR AWS ACCOUNT ID
        final String initialTopicId, // Q TOPIC ID TO WHICH THE CONSTRUCTED URL
POINTS AND SEARCHBAR PREPOPULATES INITIALLY
        final String namespace, // ANONYMOUS EMBEDDING REQUIRES SPECIFYING A
VALID NAMESPACE FOR WHICH YOU WANT THE EMBEDDING URL
        final List<String> authorizedResourceArns, // Q SEARCHBAR TOPIC ARN
LIST TO EMBED
        final List<String> allowedDomains, // RUNTIME ALLOWED DOMAINS FOR
EMBEDDING
        final List<SessionTag> sessionTags // SESSION TAGS USED FOR ROW-LEVEL
SECURITY
    ) throws Exception {
        AnonymousUserEmbeddingExperienceConfiguration experienceConfiguration =
new AnonymousUserEmbeddingExperienceConfiguration();

```



```

        AnonymousUserQSearchBarEmbeddingConfiguration qSearchBarConfiguration =
new AnonymousUserQSearchBarEmbeddingConfiguration();
        qSearchBarConfiguration.setInitialTopicId(initialTopicId);
        experienceConfiguration.setQSearchBar(qSearchBarConfiguration);

        GenerateEmbedUrlForAnonymousUserRequest
generateEmbedUrlForAnonymousUserRequest = new
GenerateEmbedUrlForAnonymousUserRequest()
                .withAwsAccountId(accountId)
                .withNamespace(namespace)
                .withAuthorizedResourceArns(authorizedResourceArns)
                .withExperienceConfiguration(experienceConfiguration)
                .withSessionTags(sessionTags)
                .withSessionLifetimeInMinutes(600L); // OPTIONAL: VALUE CAN BE
[15-600]. DEFAULT: 600
                .withAllowedDomains(allowedDomains);

        GenerateEmbedUrlForAnonymousUserResult qSearchBarEmbedUrl =
quickSightClient.generateEmbedUrlForAnonymousUser(generateEmbedUrlForAnonymousUserRequest);

        return qSearchBarEmbedUrl.getEmbedUrl();
    }
}

```

JavaScript

```

global.fetch = require('node-fetch');
const AWS = require('aws-sdk');

function generateEmbedUrlForAnonymousUser(
    accountId, // YOUR AWS ACCOUNT ID
    initialTopicId, // Q TOPIC ID TO WHICH THE CONSTRUCTED URL POINTS
    quicksightNamespace, // VALID NAMESPACE WHERE YOU WANT TO DO NOAUTH EMBEDDING
    authorizedResourceArns, // Q SEARCHBAR TOPIC ARN LIST TO EMBED
    allowedDomains, // RUNTIME ALLOWED DOMAINS FOR EMBEDDING
    sessionTags, // SESSION TAGS USED FOR ROW-LEVEL SECURITY
    generateEmbedUrlForAnonymousUserCallback, // SUCCESS CALLBACK METHOD
    errorCallback // ERROR CALLBACK METHOD
) {
    const experienceConfiguration = {
        "QSearchBar": {

```

```

        "InitialTopicId": initialTopicId // TOPIC ID CAN BE FOUND IN THE URL ON THE
TOPIC AUTHOR PAGE
    }
};

const generateEmbedUrlForAnonymousUserParams = {
    "AwsAccountId": accountId,
    "Namespace": quicksightNamespace,
    "AuthorizedResourceArns": authorizedResourceArns,
    "AllowedDomains": allowedDomains,
    "ExperienceConfiguration": experienceConfiguration,
    "SessionTags": sessionTags,
    "SessionLifetimeInMinutes": 600
};

const quicksightClient = new AWS.QuickSight({
    region: process.env.AWS_REGION,
    credentials: {
        accessKeyId: AccessKeyId,
        secretAccessKey: SecretAccessKey,
        sessionToken: SessionToken,
        expiration: Expiration
    }
});

quicksightClient.generateEmbedUrlForAnonymousUser(generateEmbedUrlForAnonymousUserParams,
function(err, data) {
    if (err) {
        console.log(err, err.stack);
        errorCallback(err);
    } else {
        const result = {
            "statusCode": 200,
            "headers": {
                "Access-Control-Allow-Origin": "*", // USE YOUR WEBSITE DOMAIN TO
SECURE ACCESS TO THIS API
                "Access-Control-Allow-Headers": "Content-Type"
            },
            "body": JSON.stringify(data),
            "isBase64Encoded": false
        }
        generateEmbedUrlForAnonymousUserCallback(result);
    }
}

```

```
});
}
```

Python3

```
import json
import boto3
from botocore.exceptions import ClientError
import time

# Create QuickSight and STS clients
quicksightClient = boto3.client('quicksight', region_name='us-west-2')
sts = boto3.client('sts')

# Function to generate embedded URL for anonymous user
# accountId: YOUR AWS ACCOUNT ID
# quicksightNamespace: VALID NAMESPACE WHERE YOU WANT TO DO NOAUTH EMBEDDING
# authorizedResourceArns: TOPIC ARN LIST TO EMBED
# allowedDomains: RUNTIME ALLOWED DOMAINS FOR EMBEDDING
# experienceConfiguration: configuration which specifies the TOPIC ID to point URL to
# sessionTags: SESSION TAGS USED FOR ROW-LEVEL SECURITY
def generateEmbedUrlForAnonymousUser(accountId, quicksightNamespace,
    authorizedResourceArns, allowedDomains, experienceConfiguration, sessionTags):
    try:
        response = quicksightClient.generate_embed_url_for_anonymous_user(
            AwsAccountId = accountId,
            Namespace = quicksightNamespace,
            AuthorizedResourceArns = authorizedResourceArns,
            AllowedDomains = allowedDomains,
            ExperienceConfiguration = experienceConfiguration,
            SessionTags = sessionTags,
            SessionLifetimeInMinutes = 600
        )

        return {
            'statusCode': 200,
            'headers': {"Access-Control-Allow-Origin": "*", "Access-Control-Allow-Headers": "Content-Type"},
            'body': json.dumps(response),
            'isBase64Encoded': bool('false')
        }
    except ClientError as e:
        print(e)
```

```
return "Error generating embeddedURL: " + str(e)
```

Node.js

The following example shows the JavaScript (Node.js) that you can use on the app server to generate the URL for the embedded dashboard. You can use this URL in your website or app to display the dashboard.

Example

```
const AWS = require('aws-sdk');
const https = require('https');

var quicksightClient = new AWS.Service({
  apiConfig: require('./quicksight-2018-04-01.min.json'),
  region: 'us-east-1',
});

quicksightClient.generateEmbedUrlForAnonymousUser({
  'AwsAccountId': '111122223333',
  'Namespace': 'DEFAULT',
  'AuthorizedResourceArns': '["topic-arn-topicId1","topic-arn-topicId2"]',
  'AllowedDomains': allowedDomains,
  'ExperienceConfiguration': {
    'QSearchBar': {
      'InitialTopicId': 'U4zJMVZ2n2stZflc80u3iKySEb3BEV6f'
    }
  },
  'SessionTags': '["Key": tag-key-1,"Value": tag-value-1,{"Key": tag-key-1,"Value": tag-value-1}]',
  'SessionLifetimeInMinutes': 15
}, function(err, data) {
  console.log('Errors: ');
  console.log(err);
  console.log('Response: ');
  console.log(data);
});
```

Example

```
//The URL returned is over 900 characters. For this example, we've shortened the string
//for
//readability and added ellipsis to indicate that it's incomplete.
```

```
{
  Status: 200,
  EmbedUrl : 'https://quicksightdomain/embed/12345/dashboards/67890/sheets/12345/
visuals/67890...',
  RequestId: '7bee030e-f191-45c4-97fe-d9faf0e03713'
}
```

.NET/C#

The following example shows the .NET/C# code that you can use on the app server to generate the URL for the embedded Q search bar. You can use this URL in your website or app to display the Q search bar.

Example

```
using System;
using Amazon.QuickSight;
using Amazon.QuickSight.Model;

namespace GenerateQSearchBarEmbedUrlForAnonymousUser
{
    class Program
    {
        static void Main(string[] args)
        {
            var quicksightClient = new AmazonQuickSightClient(
                AccessKey,
                SecretAccessKey,
                SessionToken,
                Amazon.RegionEndpoint.USEast1);
            try
            {
                AnonymousUserQSearchBarEmbeddingConfiguration
anonymousUserQSearchBarEmbeddingConfiguration
                = new AnonymousUserQSearchBarEmbeddingConfiguration
                {
                    InitialTopicId = "U4zJMVZ2n2stZflc80u3iKySEb3BEV6f"
                };
                AnonymousUserEmbeddingExperienceConfiguration
anonymousUserEmbeddingExperienceConfiguration
                = new AnonymousUserEmbeddingExperienceConfiguration
                {
                    QSearchBar = anonymousUserQSearchBarEmbeddingConfiguration

```

```
};

Console.WriteLine(
    quicksightClient.GenerateEmbedUrlForAnonymousUserAsync(new
GenerateEmbedUrlForAnonymousUserRequest
    {
        AwsAccountId = "111122223333",
        Namespace = "DEFAULT",
        AuthorizedResourceArns ["topic-arn-topicId1","topic-arn-
topicId2"]',
        AllowedDomains = allowedDomains,
        ExperienceConfiguration =
anonymousUserEmbeddingExperienceConfiguration,
        SessionTags = '[{"Key": tag-key-1,"Value": tag-value-1,{"Key":
tag-key-1,"Value": tag-value-1}}]',
        SessionLifetimeInMinutes = 15,
    }).Result.EmbedUrl
);
} catch (Exception ex) {
    Console.WriteLine(ex.Message);
}
}
```

AWS CLI

To assume the role, choose one of the following AWS Security Token Service (AWS STS) API operations:

- [AssumeRole](#) – Use this operation when you are using an IAM identity to assume the role.
- [AssumeRoleWithWebIdentity](#) – Use this operation when you are using a web identity provider to authenticate your user.
- [AssumeRoleWithSaml](#) – Use this operation when you are using SAML to authenticate your users.

The following example shows the CLI command to set the IAM role. The role needs to have permissions enabled for `quicksight:GenerateEmbedUrlForAnonymousUser`.

```
aws sts assume-role \
    --role-arn "arn:aws:iam::111122223333:role/embedding_quicksight_q_search_bar_role" \
```

```
--role-session-name anonymous caller
```

The `assume-role` operation returns three output parameters: the access key, the secret key, and the session token.

 Note

If you get an `ExpiredToken` error when calling the `AssumeRole` operation, this is probably because the previous `SESSION TOKEN` is still in the environment variables. Clear this by setting the following variables:

- `AWS_ACCESS_KEY_ID`
- `AWS_SECRET_ACCESS_KEY`
- `AWS_SESSION_TOKEN`

The following example shows how to set these three parameters in the CLI. For a Microsoft Windows machine, use `set` instead of `export`.

```
export AWS_ACCESS_KEY_ID      = "access_key_from_assume_role"
export AWS_SECRET_ACCESS_KEY  = "secret_key_from_assume_role"
export AWS_SESSION_TOKEN      = "session_token_from_assume_role"
```

Running these commands sets the role session ID of the user visiting your website to `embedding_quicksight_q_search_bar_role/QuickSightEmbeddingAnonymousPolicy`. The role session ID is made up of the role name from `role-arn` and the `role-session-name` value. Using the unique role session ID for each user ensures that appropriate permissions are set for each user. It also prevents any throttling of user access. *Throttling* is a security feature that prevents the same user from accessing Amazon Quick Sight from multiple locations. In addition, it keeps each session separate and distinct. If you're using an array of web servers, for example for load balancing, and a session is reconnected to a different server, a new session begins.

To get a signed URL for the dashboard, call `generate-embed-url-for-anonymous-user` from the app server. This returns the embeddable dashboard URL. The following example shows how to generate the URL for an embedded dashboard using a server-side call for users who are making anonymous visits to your web portal or app.

```
aws quicksight generate-embed-url-for-anonymous-user \  
--aws-account-id 111122223333 \  

```

```
--namespace default-or-something-else \  
--authorized-resource-arns '["topic-arn-topicId1","topic-arn-topicId2"]' \  
--allowed-domains '["domain1","domain2"]' \  
--experience-configuration 'QSearchBar={InitialTopicId="topicId1"}' \  
--session-tags '["Key": tag-key-1, "Value": tag-value-1, {"Key": tag-key-1, "Value": tag-value-1}]' \  
--session-lifetime-in-minutes 15
```

For more information about using this operation, see [GenerateEmbedUrlForRegisteredUser](#). You can use this and other API operations in your own code.

Step 3: Embed the Q search bar URL

Note

The embedded Amazon Quick Sight Q search bar provides the classic Amazon Quick Sight Q&A experience. Amazon Quick Sight integrates with Amazon Q Business to launch a new Generative Q&A experience. Developers are recommended to use the new Generative Q&A experience. For more information on the embedded Generative Q&A experience, see [Embedding the Amazon Q in Amazon Quick Sight Generative Q&A experience](#).

In the following section, you can find how to embed the Q search bar URL from step 3 in your website or application page. You do this with the [Amazon Quick Sight embedding SDK](#) (JavaScript). With the SDK, you can do the following:

- Place the Q search bar on an HTML page.
- Pass parameters into the Q search bar.
- Handle error states with messages that are customized to your application.

To generate the URL that you can embed in your app, call the `GenerateEmbedUrlForAnonymousUser` API operation. This URL is valid for 5 minutes, and the resulting session is valid for up to 10 hours. The API operation provides the URL with an `auth_code` value that enables a single-sign on session.

The following shows an example response from `generate-embed-url-for-anonymous-user`.

```
//The URL returned is over 900 characters. For this example, we've shortened the string  
for
```



```
//readability and added ellipsis to indicate that it's incomplete.
{
  "Status": "200",
  "EmbedUrl": "https://quicksightdomain/embedding/12345/q/search...",
  "RequestId": "7bee030e-f191-45c4-97fe-d9faf0e03713"
}
```

Embed the Q search bar in your webpage by using the [Amazon Quick Sight embedding SDK](#) or by adding this URL into an iframe. If you set a fixed height and width number (in pixels), Amazon Quick Sight uses those and doesn't change your visual as your window resizes. If you set a relative percent height and width, Amazon Quick Sight provides a responsive layout that is modified as your window size changes.

To do this, make sure that the domain to host the embedded Q search bar is on the *allow list*, the list of approved domains for your Amazon Quick Sight subscription. This requirement protects your data by keeping unapproved domains from hosting embedded Q search bar. For more information about adding domains for an embedded Q search bar, see [Managing domains and embedding](#).

When you use the Amazon Quick Sight Embedding SDK, the Q search bar on your page is dynamically resized based on the state. By using the Amazon Quick Sight Embedding SDK, you can also control parameters within the Q search bar and receive callbacks in terms of page load completion and errors.

The following example shows how to use the generated URL. This code is generated on your app server.

SDK 2.0

```
<!DOCTYPE html>
<html>

  <head>
    <title>Q Search Bar Embedding Example</title>
    <script src="https://unpkg.com/amazon-quicksight-embedding-sdk@2.0.0/dist/
quicksight-embedding-js-sdk.min.js"></script>
    <script type="text/javascript">
      const embedQSearchBar = async() => {
        const {
          createEmbeddingContext,
        } = QuickSightEmbedding;

        const embeddingContext = await createEmbeddingContext({
```

```

        onChange: (changeEvent, metadata) => {
            console.log('Context received a change', changeEvent,
metadata);
        },
    });

    const frameOptions = {
        url: "<YOUR_EMBED_URL>", // replace this value with the url
generated via embedding API
        container: '#experience-container',
        height: "700px",
        width: "1000px",
        onChange: (changeEvent, metadata) => {
            switch (changeEvent.eventName) {
                case 'FRAME_MOUNTED': {
                    console.log("Do something when the experience frame is
mounted.");

                    break;
                }
                case 'FRAME_LOADED': {
                    console.log("Do something when the experience frame is
loaded.");

                    break;
                }
            }
        },
    };

    const contentOptions = {
        hideTopicName: false,
        theme: '<YOUR_THEME_ID>',
        allowTopicSelection: true,
        onMessage: async (messageEvent, experienceMetadata) => {
            switch (messageEvent.eventName) {
                case 'Q_SEARCH_OPENED': {
                    console.log("Do something when Q Search content
expanded");

                    break;
                }
                case 'Q_SEARCH_CLOSED': {
                    console.log("Do something when Q Search content
collapsed");

                    break;
                }
            }
        }
    };

```

```

        case 'Q_SEARCH_SIZE_CHANGED': {
            console.log("Do something when Q Search size changed");
            break;
        }
        case 'CONTENT_LOADED': {
            console.log("Do something when the Q Search is
loaded.");
            break;
        }
        case 'ERROR_OCCURRED': {
            console.log("Do something when the Q Search fails
loading.");
            break;
        }
    }
};
const embeddedDashboardExperience = await
embeddingContext.embedQSearchBar(frameOptions, contentOptions);
    };
</script>
</head>

<body onload="embedQSearchBar()">
    <div id="experience-container"></div>
</body>

</html>

```

SDK 1.0

```

<!DOCTYPE html>
<html>

    <head>
        <title>QuickSight Q Search Bar Embedding</title>
        <script src="https://unpkg.com/amazon-quicksight-embedding-sdk@1.18.0/dist/
quicksight-embedding-js-sdk.min.js"></script>
        <script type="text/javascript">
            var session

            function onError(payload) {
                console.log("Do something when the session fails loading");
            }
        </script>
    </head>
</html>

```

```
    }

    function onOpen() {
        console.log("Do something when the Q search bar opens");
    }

    function onClose() {
        console.log("Do something when the Q search bar closes");
    }

    function embedQSearchBar() {
        var containerDiv = document.getElementById("embeddingContainer");
        var options = {
            url: "https://us-east-1.quicksight.aws.amazon.com/sn/dashboards/
dashboardId?isauthcode=true&identityprovider=quicksight&code=authcode", // replace this
dummy url with the one generated via embedding API
            container: containerDiv,
            width: "1000px",
            locale: "en-US",
            qSearchBarOptions: {
                expandCallback: onOpen,
                collapseCallback: onClose,
                iconDisabled: false,
                topicNameDisabled: false,
                themeId: 'bdb844d0-0fe9-4d9d-b520-0fe602d93639',
                allowTopicSelection: true
            }
        };
        session = QuickSightEmbedding.embedQSearchBar(options);
        session.on("error", onError);
    }

    function onCountryChange(obj) {
        session.setParameters({country: obj.value});
    }
</script>
</head>

<body onload="embedQSearchBar()">
    <div id="embeddingContainer"></div>
</body>

</html>
```

For this example to work, make sure to use the Amazon Quick Sight Embedding SDK to load the embedded Q search bar on your website using JavaScript. To get your copy, do one of the following:

- Download the [Amazon Quick Sight embedding SDK](#) from GitHub. This repository is maintained by a group of Amazon Quick Sight developers.
- Download the latest embedding SDK version from <https://www.npmjs.com/package/amazon-quicksight-embedding-sdk>.
- If you use npm for JavaScript dependencies, download and install it by running the following command.

```
npm install amazon-quicksight-embedding-sdk
```

Optional Amazon Quick Sight Q search bar embedding functionalities

Note

The embedded Amazon Quick Sight Q search bar provides the classic Amazon Quick Sight Q&A experience. Amazon Quick Sight integrates with Amazon Q Business to launch a new Generative Q&A experience. Developers are recommended to use the new Generative Q&A experience. For more information on the embedded Generative Q&A experience, see [Embedding the Amazon Q in Amazon Quick Sight Generative Q&A experience](#).

The following optional functionalities are available for the embedded Q search bar using the embedding SDK.

Invoke Q search bar actions

The following options are only supported for Q search bar embedding.

- Set a Q search bar question — This feature sends a question to the Q search bar and immediately queries the question. It also automatically opens the Q popover.

```
qBar.setQBarQuestion('show me monthly revenue');
```

- Close the Q popover — This feature closes the Q popover and returns the iframe to the original Q search bar size.

```
qBar.closeQPopover();
```

For more information, see the [Amazon Quick Sight embedding SDK](#).

Embedding analytics using the `GetDashboardEmbedURL` and `GetSessionEmbedURL` API operations

Applies to: Enterprise Edition

Intended audience: Amazon Quick developers

The following API operations for embedding Amazon Quick Sight dashboards and the Amazon Quick Sight console have been replaced by the `GenerateEmbedUrlForAnonymousUser` and `GenerateEmbedUrlForRegisteredUser` API operations. You can still use them to embed analytics in your application, but they are no longer maintained and do not contain the latest embedding features or functionality. For the latest up-to-date embedding experience, see [Embedding Amazon Quick Sight analytics into your applications](#)

- The [GetDashboardEmbedUrl](#) API operation embeds interactive dashboards.
- The [GetSessionEmbedUrl](#) API operation embeds the Amazon Quick Sight console.

Topics

- [Embedding dashboards for everyone using `GetDashboardEmbedURL` \(old API\)](#)
- [Embedding dashboards for registered users using `GetDashboardEmbedUrl` \(old API\)](#)
- [Embedding the Amazon Quick Sight console using `GetSessionEmbedUrl` \(old API\)](#)

Embedding dashboards for everyone using `GetDashboardEmbedURL` (old API)

Important

Amazon Quick Sight has new APIs for embedding analytics:
`GenerateEmbedUrlForAnonymousUser` and `GenerateEmbedUrlForRegisteredUser`.

You can still use the `GetDashboardEmbedUrl` and `GetSessionEmbedUrl` APIs to embed dashboards and the Amazon Quick Sight console, but they do not contain the latest embedding capabilities. For the latest up-to-date embedding experience, see [Embedding Amazon Quick Sight analytics into your applications](#).

Applies to: Enterprise Edition

Intended audience: Amazon Quick developers

In the following sections, you can find detailed information on how to set up embedded Amazon Quick Sight dashboards for everyone (nonauthenticated users) using `GetDashboardEmbedURL`.

Topics

- [Step 1: Set up permissions](#)
- [Step 2: Get the URL with the authentication code attached](#)
- [Step 3: Embed the dashboard URL](#)

Step 1: Set up permissions

Important

Amazon Quick Sight has new APIs for embedding analytics: `GenerateEmbedUrlForAnonymousUser` and `GenerateEmbedUrlForRegisteredUser`. You can still use the `GetDashboardEmbedUrl` and `GetSessionEmbedUrl` APIs to embed dashboards and the Amazon Quick Sight console, but they do not contain the latest embedding capabilities. For the latest up-to-date embedding experience, see [Embedding Amazon Quick Sight analytics into your applications](#).

Applies to: Enterprise Edition

Intended audience: Amazon Quick developers

In the following section, you can find out how to set up permissions for the backend application or web server. This task requires administrative access to IAM.

Each user who accesses a dashboard assumes a role that gives them Amazon Quick Sight access and permissions to the dashboard. To make this possible, create an IAM role in your AWS account. Associate an IAM policy with the role to provide permissions to any user who assumes it.

The following sample policy provides these permissions for use with `IdentityType=ANONYMOUS`. For this approach to work, you also need a session pack, or session capacity pricing, on your AWS account. Otherwise, when a user tries to access the dashboard, the error `UnsupportedPricingPlanException` is returned.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "quicksight:GetDashboardEmbedUrl",
        "quicksight:GetAnonymousUserEmbedUrl"
      ],
      "Resource": "*"
    }
  ]
}
```

Your application's IAM identity must have a trust policy associated with it to allow access to the role that you just created. This means that when a user accesses your application, your application can assume the role on the user's behalf to open the dashboard. The following example shows a role called `QuickSightEmbeddingAnonymousPolicy`, which has the sample policy preceding as its resource.

For more information regarding trust policies, see [Temporary security credentials in IAM](#) in the *IAM User Guide*.

Step 2: Get the URL with the authentication code attached

Important

Amazon Quick Sight has new APIs for embedding analytics: `GenerateEmbedUrlForAnonymousUser` and `GenerateEmbedUrlForRegisteredUser`. You can still use the `GetDashboardEmbedUrl` and `GetSessionEmbedUrl` APIs to embed dashboards and the Amazon Quick Sight console, but they do not contain the latest embedding capabilities. For the latest up-to-date embedding experience, see [Embedding Amazon Quick Sight analytics into your applications](#).

Applies to: Enterprise Edition

Intended audience: Amazon Quick developers

In the following section, you can find how to authenticate on behalf of the anonymous visitor and get the embeddable dashboard URL on your application server.

When a user accesses your app, the app assumes the IAM role on the user's behalf. Then it adds the user to Amazon Quick Sight, if that user doesn't already exist. Next, it passes an identifier as the unique role session ID.

The following examples perform the IAM authentication on the user's behalf. It passes an identifier as the unique role session ID. This code runs on your app server.

Java

```
import com.amazonaws.auth.AWSCredentials;
import com.amazonaws.auth.BasicAWSCredentials;
import com.amazonaws.auth.AWSCredentialsProvider;
import com.amazonaws.regions.Regions;
import com.amazonaws.services.quicksight.AmazonQuickSight;
import com.amazonaws.services.quicksight.AmazonQuickSightClientBuilder;
import com.amazonaws.services.quicksight.model.GetDashboardEmbedUrlRequest;
import com.amazonaws.services.quicksight.model.GetDashboardEmbedUrlResult;
```

```

/**
 * Class to call QuickSight AWS SDK to get url for dashboard embedding.
 */
public class GetQuicksightEmbedUrlNoAuth {

    private static String ANONYMOUS = "ANONYMOUS";

    private final AmazonQuickSight quickSightClient;

    public GetQuicksightEmbedUrlNoAuth() {
        this.quickSightClient = AmazonQuickSightClientBuilder
            .standard()
            .withRegion(Regions.US_EAST_1.getName())
            .withCredentials(new AWSCredentialsProvider() {
                @Override
                public AWSCredentials getCredentials() {
                    // provide actual IAM access key and secret
                    key here
                    return new BasicAWSCredentials("access-
key", "secret-key");
                }

                @Override
                public void refresh() {}
            })
            .build();
    }

    public String getQuicksightEmbedUrl(
        final String accountId, // YOUR AWS ACCOUNT ID
        final String dashboardId, // YOUR DASHBOARD ID TO EMBED
        final String additionalDashboardIds, // ADDITIONAL DASHBOARD-1 ADDITIONAL
        DASHBOARD-2
        final boolean resetDisabled, // OPTIONAL PARAMETER TO ENABLE DISABLE
        RESET BUTTON IN EMBEDDED DASHBAORD
        final boolean undoRedoDisabled // OPTIONAL PARAMETER TO ENABLE DISABLE
        UNDO REDO BUTTONS IN EMBEDDED DASHBAORD
    ) throws Exception {
        GetDashboardEmbedUrlRequest getDashboardEmbedUrlRequest = new
        GetDashboardEmbedUrlRequest()
            .withDashboardId(dashboardId)
            .withAdditionalDashboardIds(additionalDashboardIds)
            .withAwsAccountId(accountId)
    }

```

```

        .withNamespace("default") // Anonymous embedding requires specifying
        a valid namespace for which you want the embedding url
        .withIdentityType(ANONYMOUS)
        .withResetDisabled(resetDisabled)
        .withUndoRedoDisabled(undoRedoDisabled);

        GetDashboardEmbedUrlResult dashboardEmbedUrl =
        quickSightClient.getDashboardEmbedUrl(getDashboardEmbedUrlRequest);

        return dashboardEmbedUrl.getEmbedUrl();
    }
}

```

JavaScript

```

global.fetch = require('node-fetch');
const AWS = require('aws-sdk');

function getDashboardEmbedURL(
    accountId, // YOUR AWS ACCOUNT ID
    dashboardId, // YOUR DASHBOARD ID TO EMBED
    additionalDashboardIds, // ADDITIONAL DASHBOARD-1 ADDITIONAL DASHBOARD-2
    quicksightNamespace, // VALID NAMESPACE WHERE YOU WANT TO DO NOAUTH EMBEDDING
    resetDisabled, // OPTIONAL PARAMETER TO ENABLE DISABLE RESET BUTTON IN EMBEDDED
    DASHBAORD
    undoRedoDisabled, // OPTIONAL PARAMETER TO ENABLE DISABLE UNDO REDO BUTTONS IN
    EMBEDDED DASHBAORD
    getEmbedUrlCallback, // GETEMBEDURL SUCCESS CALLBACK METHOD
    errorCallback // GETEMBEDURL ERROR CALLBACK METHOD
) {
    const getDashboardParams = {
        AwsAccountId: accountId,
        DashboardId: dashboardId,
        AdditionalDashboardIds: additionalDashboardIds,
        Namespace: quicksightNamespace,
        IdentityType: 'ANONYMOUS',
        ResetDisabled: resetDisabled,
        SessionLifetimeInMinutes: 600,
        UndoRedoDisabled: undoRedoDisabled
    };

    const quicksightGetDashboard = new AWS.QuickSight({
        region: process.env.AWS_REGION,

```

```

    });

    quicksightGetDashboard.getDashboardEmbedUrl(getDashboardParams, function(err,
data) {
    if (err) {
        console.log(err, err.stack);
        errorCallback(err);
    } else {
        const result = {
            "statusCode": 200,
            "headers": {
                "Access-Control-Allow-Origin": "*", // USE YOUR WEBSITE DOMAIN
                "Access-Control-Allow-Headers": "Content-Type"
            },
            "body": JSON.stringify(data),
            "isBase64Encoded": false
        }
        getEmbedUrlCallback(result);
    }
    });
}

```

Python3

```

import json
import boto3
from botocore.exceptions import ClientError
import time

# Create QuickSight and STS clients
qs = boto3.client('quicksight', region_name='us-east-1')
sts = boto3.client('sts')

# Function to generate embedded URL
# accountId: YOUR AWS ACCOUNT ID
# dashboardId: YOUR DASHBOARD ID TO EMBED
# additionalDashboardIds: ADDITIONAL DASHBOARD-1 ADDITIONAL DASHBOARD-2 WITHOUT
# COMMAS
# quicksightNamespace: VALID NAMESPACE WHERE YOU WANT TO DO NOAUTH EMBEDDING
# resetDisabled: PARAMETER TO ENABLE DISABLE RESET BUTTON IN EMBEDDED DASHBAORD
# undoRedoDisabled: OPTIONAL PARAMETER TO ENABLE DISABLE UNDO REDO BUTTONS IN
# EMBEDDED DASHBAORD

```

```
def getDashboardURL(accountId, dashboardId, quicksightNamespace, resetDisabled,
undoRedoDisabled):
    try:
        response = qs.get_dashboard_embed_url(
            AwsAccountId = accountId,
            DashboardId = dashboardId,
            AdditionalDashboardIds = additionalDashboardIds,
            Namespace = quicksightNamespace,
            IdentityType = 'ANONYMOUS',
            SessionLifetimeInMinutes = 600,
            UndoRedoDisabled = undoRedoDisabled,
            ResetDisabled = resetDisabled
        )

        return {
            'statusCode': 200,
            'headers': {"Access-Control-Allow-Origin": "*", "Access-Control-Allow-
Headers": "Content-Type"},
            'body': json.dumps(response),
            'isBase64Encoded': bool('false')
        }
    except ClientError as e:
        print(e)
        return "Error generating embeddedURL: " + str(e)
```

Node.js

The following example shows the JavaScript (Node.js) that you can use on the app server to get the URL for the embedded dashboard. You can use this URL in your website or app to display the dashboard.

Example

```
const AWS = require('aws-sdk');
const https = require('https');

var quicksight = new AWS.Service({
    apiConfig: require('./quicksight-2018-04-01.min.json'),
    region: 'us-east-1',
});

quicksight.getDashboardEmbedUrl({
    'AwsAccountId': '111122223333',
```

```

        'DashboardId': 'dashboard-id',
        'AdditionalDashboardIds': 'added-dashboard-id-1 added-dashboard-id-2
added-dashboard-id-3'
        'Namespace' : 'default',
        'IdentityType': 'ANONYMOUS',
        'SessionLifetimeInMinutes': 100,
        'UndoRedoDisabled': false,
        'ResetDisabled': true

    }, function(err, data) {
        console.log('Errors: ');
        console.log(err);
        console.log('Response: ');
        console.log(data);
    });

```

Example

```

//The URL returned is over 900 characters. For this example, we've shortened the
string for
    //readability and added ellipsis to indicate that it's incomplete.
        { Status: 200,
          EmbedUrl: 'https://dashboards.example.com/
embed/620bef10822743fab329fb3751187d2d...
          RequestId: '7bee030e-f191-45c4-97fe-d9faf0e03713' }

```

.NET/C#

The following example shows the .NET/C# code that you can use on the app server to get the URL for the embedded dashboard. You can use this URL in your website or app to display the dashboard.

Example

```

var client = new AmazonQuickSightClient(
    AccessKey,
    SecretAccessKey,
    sessionToken,
    Amazon.RegionEndpoint.USEast1);
try

```

```

    {
        Console.WriteLine(
            client.GetDashboardEmbedUrlAsync(new GetDashboardEmbedUrlRequest
            {
                AwsAccountId = "111122223333",
                DashboardId = "dashboard-id",
                AdditionalDashboardIds = "added-dashboard-id-1 added-
dashboard-id-2 added-dashboard-id-3",
                Namespace = default,
                IdentityType = IdentityType.ANONYMOUS,
                SessionLifetimeInMinutes = 600,
                UndoRedoDisabled = false,
                ResetDisabled = true
            }).Result.EmbedUrl
        );
    } catch (Exception ex) {
        Console.WriteLine(ex.Message);
    }
}

```

AWS CLI

To assume the role, choose one of the following AWS Security Token Service (AWS STS) API operations:

- [AssumeRole](#) – Use this operation when you are using an IAM identity to assume the role.
- [AssumeRoleWithWebIdentity](#) – Use this operation when you are using a web identity provider to authenticate your user.
- [AssumeRoleWithSaml](#) – Use this operation when you are using Security Assertion Markup Language (SAML) to authenticate your users.

The following example shows the CLI command to set the IAM role. The role needs to have permissions enabled for `quicksight:GetDashboardEmbedURL`.

```

aws sts assume-role \
    --role-arn "arn:aws:iam::111122223333:role/QuickSightEmbeddingAnonymousPolicy" \
    --role-session-name anonymous caller

```

The `assume-role` operation returns three output parameters: the access key, the secret key, and the session token.

Note

If you get an ExpiredToken error when calling the AssumeRole operation, this is probably because the previous SESSION_TOKEN is still in the environment variables. Clear this by setting the following variables:

- `AWS_ACCESS_KEY_ID`
- `AWS_SECRET_ACCESS_KEY`
- `AWS_SESSION_TOKEN`

The following example shows how to set these three parameters in the CLI. If you are using a Microsoft Windows machine, use `set` instead of `export`.

```
export AWS_ACCESS_KEY_ID      = "access_key_from_assume_role"
export AWS_SECRET_ACCESS_KEY  = "secret_key_from_assume_role"
export AWS_SESSION_TOKEN      = "session_token_from_assume_role"
```

Running these commands sets the role session ID of the user visiting your website to `embedding_quicksight_dashboard_role/QuickSightEmbeddingAnonymousPolicy`. The role session ID is made up of the role name from `role-arn` and the `role-session-name` value. Using the unique role session ID for each user ensures that appropriate permissions are set for each visiting user. It also keeps each session separate and distinct. If you're using an array of web servers, for example for load balancing, and a session is reconnected to a different server, a new session begins.

To get a signed URL for the dashboard, call `get-dashboard-embed-url` from the app server. This returns the embeddable dashboard URL. The following example shows how to get the URL for an embedded dashboard using a server-side call for users who are making anonymous visits to your web portal or app.

```
aws quicksight get-dashboard-embed-url \
  --aws-account-id 111122223333 \
  --dashboard-id dashboard-id \
  --additional-dashboard-ids added-dashboard-id-1 added-dashboard-id-2 added-
  dashboard-id-3
  --namespace default-or-something-else \
  --identity-type ANONYMOUS \
  --session-lifetime-in-minutes 30 \
```



```
--undo-redo-disabled true \  
--reset-disabled true \  
--user-arn arn:aws:quicksight:us-east-1:111122223333:user/  
default/QuickSightEmbeddingAnonymousPolicy/embeddingSession
```

For more information on using this operation, see [GetDashboardEmbedUrl](#). You can use this and other API operations in your own code.

Step 3: Embed the dashboard URL

Important

Amazon Quick Sight has new APIs for embedding analytics:

`GenerateEmbedUrlForAnonymousUser` and `GenerateEmbedUrlForRegisteredUser`. You can still use the `GetDashboardEmbedUrl` and `GetSessionEmbedUrl` APIs to embed dashboards and the Amazon Quick Sight console, but they do not contain the latest embedding capabilities. For the latest up-to-date embedding experience, see [Embedding Amazon Quick Sight analytics into your applications](#).

Applies to: Enterprise Edition

Intended audience: Amazon Quick developers

In the following section, you can find out how you can use the [Amazon Quick Sight embedding SDK](#) (JavaScript) to embed the dashboard URL from step 2 in your website or application page. With the SDK, you can do the following:

- Place the dashboard on an HTML page.
- Pass parameters into the dashboard.
- Handle error states with messages that are customized to your application.

Call the `GetDashboardEmbedUrl` API operation to get the URL that you can embed in your app. This URL is valid for 5 minutes, and the resulting session is valid for 10 hours. The API operation provides the URL with an `auth_code` that enables a single-sign on session.

The following shows an example response from `get-dashboard-embed-url`.

```
//The URL returned is over 900 characters. For this example, we've shortened the string
//for
//readability and added ellipsis to indicate that it's incomplete.
{
  "Status": "200",
  "EmbedUrl": "https://dashboards.example.com/
embed/620bef10822743fab329fb3751187d2d...",
  "RequestId": "7bee030e-f191-45c4-97fe-d9faf0e03713"
}
```

Embed this dashboard in your web page by using the Amazon Quick Sight [Embedding SDK](#) or by adding this URL into an `iframe`. If you set a fixed height and width number (in pixels), Amazon Quick Sight uses those and doesn't change your visual as your window resizes. If you set a relative percent height and width, Amazon Quick Sight provides a responsive layout that is modified as your window size changes. By using the Amazon Quick Sight Embedding SDK, you can also control parameters within the dashboard and receive callbacks in terms of page load completion and errors.

The following example shows how to use the generated URL. This code resides on your app server.

```
<!DOCTYPE html>
<html>

<head>
  <title>Basic Embed</title>
  <!-- You can download the latest QuickSight embedding SDK version from https://
www.npmjs.com/package/amazon-quicksight-embedding-sdk -->
  <!-- Or you can do "npm install amazon-quicksight-embedding-sdk", if you use npm
for javascript dependencies -->
  <script src="./quicksight-embedding-js-sdk.min.js"></script>
  <script type="text/javascript">
    var dashboard;

    function embedDashboard() {
      var containerDiv = document.getElementById("embeddingContainer");
```

```
var options = {
    // replace this dummy url with the one generated via embedding API
    url: "https://us-east-1.quicksight.aws.amazon.com/sn/dashboards/
dashboardId?isauthcode=true&identityprovider=quicksight&code=authcode",
    container: containerDiv,
    scrolling: "no",
    height: "700px",
    width: "1000px",
    footerPaddingEnabled: true
};
dashboard = QuickSightEmbedding.embedDashboard(options);
}
</script>
</head>

<body onload="embedDashboard()">
    <div id="embeddingContainer"></div>
</body>

</html>
```

For this example to work, make sure to use the Amazon Quick Sight Embedding SDK to load the embedded dashboard on your website using JavaScript. To get your copy, do one of the following:

- Download the [Amazon Quick Sight embedding SDK](#) from GitHub. This repository is maintained by a group of Amazon Quick Sight developers.
- Download the latest QuickSight embedding SDK version from <https://www.npmjs.com/package/amazon-quicksight-embedding-sdk>.
- If you use npm for JavaScript dependencies, download and install it by running the following command.

```
npm install amazon-quicksight-embedding-sdk
```

Embedding dashboards for registered users using GetDashboardEmbedUrl (old API)

Important

Amazon Quick Sight has new APIs for embedding analytics:

`GenerateEmbedUrlForAnonymousUser` and `GenerateEmbedUrlForRegisteredUser`.

You can still use the `GetDashboardEmbedUrl` and `GetSessionEmbedUrl` APIs to embed dashboards and the Amazon Quick Sight console, but they do not contain the latest embedding capabilities. For the latest up-to-date embedding experience, see [Embedding Amazon Quick Sight analytics into your applications](#).

In the following sections, you can find detailed information on how to set up embedded Amazon Quick Sight dashboards for registered users using `GetDashboardEmbedUrl`.

Topics

- [Step 1: Set up permissions](#)
- [Step 2: Get the URL with the authentication code attached](#)
- [Step 3: Embed the dashboard URL](#)

Step 1: Set up permissions

Important

Amazon Quick Sight has new APIs for embedding analytics: `GenerateEmbedUrlForAnonymousUser` and `GenerateEmbedUrlForRegisteredUser`. You can still use the `GetDashboardEmbedUrl` and `GetSessionEmbedUrl` APIs to embed dashboards and the Amazon Quick Sight console, but they do not contain the latest embedding capabilities. For the latest up-to-date embedding experience, see [Embedding Amazon Quick Sight analytics into your applications](#).

In the following section, you can find out how to set up permissions for the backend application or web server. This task requires administrative access to IAM.

Each user who accesses a dashboard assumes a role that gives them Amazon Quick Sight access and permissions to the dashboard. To make this possible, create an IAM role in your AWS account. Associate an IAM policy with the role to provide permissions to any user who assumes it. The IAM role needs to provide permissions to retrieve dashboard URLs. For this, you add `quicksight:GetDashboardEmbedUrl`.

The following sample policy provides these permissions for use with `IdentityType=IAM`.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "quicksight:GetDashboardEmbedUrl"
      ],
      "Resource": "*"
    }
  ]
}
```

The following sample policy provides permission to retrieve a dashboard URL. You use the policy with `quicksight:RegisterUser` if you are creating first-time users who are to be Amazon Quick Sight readers.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": "quicksight:RegisterUser",
      "Resource": "*",
      "Effect": "Allow"
    },
    {
      "Action": "quicksight:GetDashboardEmbedUrl",
      "Resource": "*",
      "Effect": "Allow"
    }
  ]
}
```

If you use QUICKSIGHT as your `identityType` and provide the user's Amazon Resource Name (ARN), you also need to allow the `quicksight:GetAuthCode` action in your policy. The following sample policy provides this permission.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "quicksight:GetDashboardEmbedUrl",
        "quicksight:GetAuthCode"
      ],
      "Resource": "*"
    }
  ]
}
```

Your application's IAM identity must have a trust policy associated with it to allow access to the role that you just created. This means that when a user accesses your application, your application can assume the role on the user's behalf and provision the user in Amazon Quick Sight. The following example shows a role called `embedding_quicksight_dashboard_role`, which has the sample policy preceding as its resource.

For more information regarding trust policies for OpenID Connect or SAML authentication, see the following sections of the *IAM User Guide*:

- [Creating a role for web identity or OpenID Connect federation \(console\)](#)
- [Creating a role for SAML 2.0 federation \(console\)](#)

Step 2: Get the URL with the authentication code attached

Important

Amazon Quick Sight has new APIs for embedding analytics:
`GenerateEmbedUrlForAnonymousUser` and `GenerateEmbedUrlForRegisteredUser`.

You can still use the `GetDashboardEmbedUrl` and `GetSessionEmbedUrl` APIs to embed dashboards and the Amazon Quick Sight console, but they do not contain the latest embedding capabilities. For the latest up-to-date embedding experience, see [Embedding Amazon Quick Sight analytics into your applications](#).

In the following section, you can find out how to authenticate your user and get the embeddable dashboard URL on your application server.

When a user accesses your app, the app assumes the IAM role on the user's behalf. Then it adds the user to Amazon Quick Sight, if that user doesn't already exist. Next, it passes an identifier as the unique role session ID.

Performing the described steps ensures that each viewer of the dashboard is uniquely provisioned in Amazon Quick Sight. It also enforces per-user settings, such as the row-level security and dynamic defaults for parameters.

The following examples perform the IAM authentication on the user's behalf. This code runs on your app server.

Java

```
import com.amazonaws.auth.AWSCredentials;
import com.amazonaws.auth.AWSStaticCredentialsProvider;
import com.amazonaws.auth.BasicSessionCredentials;
import com.amazonaws.auth.BasicAWSCredentials;
import com.amazonaws.auth.AWSCredentialsProvider;
import com.amazonaws.regions.Regions;
import com.amazonaws.services.quicksight.AmazonQuickSight;
import com.amazonaws.services.quicksight.AmazonQuickSightClientBuilder;
import com.amazonaws.services.quicksight.model.GetDashboardEmbedUrlRequest;
import com.amazonaws.services.quicksight.model.GetDashboardEmbedUrlResult;
import com.amazonaws.services.securitytoken.AWSSecurityTokenService;
import com.amazonaws.services.securitytoken.model.AssumeRoleRequest;
import com.amazonaws.services.securitytoken.model.AssumeRoleResult;

/**
 * Class to call QuickSight AWS SDK to get url for dashboard embedding.
 */
public class GetQuicksightEmbedUrlIAMAuth {

    private static String IAM = "IAM";
```

```

private final AmazonQuickSight quickSightClient;

private final AWSSecurityTokenService awsSecurityTokenService;

public GetQuicksightEmbedUrlIAMAuth(final AWSSecurityTokenService
awsSecurityTokenService) {
    this.quickSightClient = AmazonQuickSightClientBuilder
        .standard()
        .withRegion(Regions.US_EAST_1.getName())
        .withCredentials(new AWSCredentialsProvider() {
            @Override
            public AWSCredentials getCredentials() {
                // provide actual IAM access key and secret
                return new BasicAWSCredentials("access-
key here", "secret-key");
            }

            @Override
            public void refresh() {}
        })
        .build();
    this.awsSecurityTokenService = awsSecurityTokenService;
}

public String getQuicksightEmbedUrl(
    final String accountId, // YOUR AWS ACCOUNT ID
    final String dashboardId, // YOUR DASHBOARD ID TO EMBED
    final String openIdToken, // TOKEN TO ASSUME ROLE WITH ROLEARNS
    final String roleArn, // IAM USER ROLE TO USE FOR EMBEDDING
    final String sessionName, // SESSION NAME FOR THE ROLEARNS ASSUME ROLE
    final boolean resetDisabled, // OPTIONAL PARAMETER TO ENABLE DISABLE
    RESET BUTTON IN EMBEDDED DASHBAORD
    final boolean undoRedoDisabled // OPTIONAL PARAMETER TO ENABLE DISABLE
    UNDO REDO BUTTONS IN EMBEDDED DASHBAORD
) throws Exception {
    AssumeRoleRequest request = new AssumeRoleRequest()
        .withRoleArn(roleArn)
        .withRoleSessionName(sessionName)
        .withTokenCode(openIdToken)
        .withDurationSeconds(3600);

```



```

        AssumeRoleResult assumeRoleResult =
awsSecurityTokenService.assumeRole(request);

        AWSCredentials temporaryCredentials = new BasicSessionCredentials(
            assumeRoleResult.getCredentials().getAccessKeyId(),
            assumeRoleResult.getCredentials().getSecretAccessKey(),
            assumeRoleResult.getCredentials().getSessionToken());
        AWSStaticCredentialsProvider awsStaticCredentialsProvider = new
AWSStaticCredentialsProvider(temporaryCredentials);

        GetDashboardEmbedUrlRequest getDashboardEmbedUrlRequest = new
GetDashboardEmbedUrlRequest()
            .withDashboardId(dashboardId)
            .withAwsAccountId(accountId)
            .withIdentityType(IAM)
            .withResetDisabled(resetDisabled)
            .withUndoRedoDisabled(undoRedoDisabled)
            .withRequestCredentialsProvider(awsStaticCredentialsProvider);

        GetDashboardEmbedUrlResult dashboardEmbedUrl =
quickSightClient.getDashboardEmbedUrl(getDashboardEmbedUrlRequest);

        return dashboardEmbedUrl.getEmbedUrl();
    }
}

```

JavaScript

```

global.fetch = require('node-fetch');
const AWS = require('aws-sdk');

function getDashboardEmbedURL(
    accountId, // YOUR AWS ACCOUNT ID
    dashboardId, // YOUR DASHBOARD ID TO EMBED
    openIdToken, // TOKEN TO ASSUME ROLE WITH ROLEARN
    roleArn, // IAM USER ROLE TO USE FOR EMBEDDING
    sessionName, // SESSION NAME FOR THE ROLEARN ASSUME ROLE
    resetDisabled, // OPTIONAL PARAMETER TO ENABLE DISABLE RESET BUTTON IN EMBEDDED
    DASHBAORD
    undoRedoDisabled, // OPTIONAL PARAMETER TO ENABLE DISABLE UNDO REDO BUTTONS IN
    EMBEDDED DASHBAORD
    getEmbedUrlCallback, // GETEMBEDURL SUCCESS CALLBACK METHOD
    errorCallback // GETEMBEDURL ERROR CALLBACK METHOD
) {
    // Implementation logic
}

```

```

    ) {
    const stsClient = new AWS.STS();
    let stsParams = {
      RoleSessionName: sessionName,
      WebIdentityToken: openIdToken,
      RoleArn: roleArn
    }

    stsClient.assumeRoleWithWebIdentity(stsParams, function(err, data) {
      if (err) {
        console.log('Error assuming role');
        console.log(err, err.stack);
        errorCallback(err);
      } else {
        const getDashboardParams = {
          AwsAccountId: accountId,
          DashboardId: dashboardId,
          IdentityType: 'IAM',
          ResetDisabled: resetDisabled,
          SessionLifetimeInMinutes: 600,
          UndoRedoDisabled: undoRedoDisabled
        };

        const quicksightGetDashboard = new AWS.QuickSight({
          region: process.env.AWS_REGION,
          credentials: {
            accessKeyId: data.Credentials.AccessKeyId,
            secretAccessKey: data.Credentials.SecretAccessKey,
            sessionToken: data.Credentials.SessionToken,
            expiration: data.Credentials.Expiration
          }
        });

        quicksightGetDashboard.getDashboardEmbedUrl(getDashboardParams,
function(err, data) {
      if (err) {
        console.log(err, err.stack);
        errorCallback(err);
      } else {
        const result = {
          "statusCode": 200,
          "headers": {
            "Access-Control-Allow-Origin": "*", // USE YOUR WEBSITE
DOMAIN TO SECURE ACCESS TO GETEMBEDURL API

```

```

        "Access-Control-Allow-Headers": "Content-Type"
    },
    "body": JSON.stringify(data),
    "isBase64Encoded": false
}
getEmbedUrlCallback(result);
    }
});
}
});
}
}

```

Python3

```

import json
import boto3
from botocore.exceptions import ClientError

# Create QuickSight and STS clients
qs = boto3.client('quicksight', region_name='us-east-1')
sts = boto3.client('sts')

# Function to generate embedded URL
# accountId: YOUR AWS ACCOUNT ID
# dashboardId: YOUR DASHBOARD ID TO EMBED
# openIdToken: TOKEN TO ASSUME ROLE WITH ROLEARN
# roleArn: IAM USER ROLE TO USE FOR EMBEDDING
# sessionName: SESSION NAME FOR THE ROLEARN ASSUME ROLE
# resetDisabled: PARAMETER TO ENABLE DISABLE RESET BUTTON IN EMBEDDED DASHBAORD
# undoRedoDisabled: PARAMETER TO ENABLE DISABLE UNDO REDO BUTTONS IN EMBEDDED
# DASHBAORD
def getDashboardURL(accountId, dashboardId, openIdToken, roleArn, sessionName,
    resetDisabled, undoRedoDisabled):
    try:
        assumedRole = sts.assume_role(
            RoleArn = roleArn,
            RoleSessionName = sessionName,
            WebIdentityToken = openIdToken
        )
    except ClientError as e:
        return "Error assuming role: " + str(e)
    else:
        assumedRoleSession = boto3.Session(

```

```

        aws_access_key_id = assumedRole['Credentials']['AccessKeyId'],
        aws_secret_access_key = assumedRole['Credentials']['SecretAccessKey'],
        aws_session_token = assumedRole['Credentials']['SessionToken'],
    )
    try:
        quickSight = assumedRoleSession.client('quicksight', region_name='us-
east-1')

        response = quickSight.get_dashboard_embed_url(
            AwsAccountId = accountId,
            DashboardId = dashboardId,
            IdentityType = 'IAM',
            SessionLifetimeInMinutes = 600,
            UndoRedoDisabled = undoRedoDisabled,
            ResetDisabled = resetDisabled
        )

        return {
            'statusCode': 200,
            'headers': {"Access-Control-Allow-Origin": "*", "Access-Control-
Allow-Headers": "Content-Type"},
            'body': json.dumps(response),
            'isBase64Encoded': bool('false')
        }
    except ClientError as e:
        return "Error generating embeddedURL: " + str(e)

```

Node.js

The following example shows the JavaScript (Node.js) that you can use on the app server to get the URL for the embedded dashboard. You can use this URL in your website or app to display the dashboard.

Example

```

const AWS = require('aws-sdk');
const https = require('https');

var quicksight = new AWS.Service({
    apiConfig: require('./quicksight-2018-04-01.min.json'),
    region: 'us-east-1',
});

```

```
quicksight.getDashboardEmbedUrl({
  'AwsAccountId': '111122223333',
  'DashboardId': '1c1fe111-e2d2-3b30-44ef-a0e111111cde',
  'IdentityType': 'IAM',
  'ResetDisabled': true,
  'SessionLifetimeInMinutes': 100,
  'UndoRedoDisabled': false,
  'StatePersistenceEnabled': true
}, function(err, data) {
  console.log('Errors: ');
  console.log(err);
  console.log('Response: ');
  console.log(data);
});
```

Example

```
//The URL returned is over 900 characters. For this example, we've shortened the
string for
    //readability and added ellipsis to indicate that it's incomplete.
    { Status: 200,
      EmbedUrl: 'https://dashboards.example.com/
embed/620bef10822743fab329fb3751187d2d...
      RequestId: '7bee030e-f191-45c4-97fe-d9faf0e03713' }
```

.NET/C#

The following example shows the .NET/C# code that you can use on the app server to get the URL for the embedded dashboard. You can use this URL in your website or app to display the dashboard.

Example

```
var client = new AmazonQuickSightClient(
    AccessKey,
    SecretAccessKey,
    sessionToken,
    Amazon.RegionEndpoint.USEast1);
try
{
```

```
Console.WriteLine(
    client.GetDashboardEmbedUrlAsync(new GetDashboardEmbedUrlRequest
    {
        AwsAccountId = "111122223333",
        DashboardId = "1c1fe111-e2d2-3b30-44ef-a0e111111cde",
        IdentityType = EmbeddingIdentityType.IAM,
        ResetDisabled = true,
        SessionLifetimeInMinutes = 100,
        UndoRedoDisabled = false,
        StatePersistenceEnabled = true
    }).Result.EmbedUrl
);
} catch (Exception ex) {
    Console.WriteLine(ex.Message);
}
```

AWS CLI

To assume the role, choose one of the following AWS Security Token Service (AWS STS) API operations:

- [AssumeRole](#) – Use this operation when you are using an IAM identity to assume the role.
- [AssumeRoleWithWebIdentity](#) – Use this operation when you are using a web identity provider to authenticate your user.
- [AssumeRoleWithSaml](#) – Use this operation when you are using SAML to authenticate your users.

The following example shows the CLI command to set the IAM role. The role needs to have permissions enabled for `quicksight:GetDashboardEmbedURL`. If you are taking a just-in-time approach to add users when they first open a dashboard, the role also needs permissions enabled for `quicksight:RegisterUser`.

```
aws sts assume-role \
    --role-arn "arn:aws:iam::111122223333:role/embedding_quicksight_dashboard_role" \
    --role-session-name john.doe@example.com
```

The `assume-role` operation returns three output parameters: the access key, the secret key, and the session token.

Note

If you get an `ExpiredToken` error when calling the `AssumeRole` operation, this is probably because the previous `SESSION_TOKEN` is still in the environment variables. Clear this by setting the following variables:

- `AWS_ACCESS_KEY_ID`
- `AWS_SECRET_ACCESS_KEY`
- `AWS_SESSION_TOKEN`

The following example shows how to set these three parameters in the CLI. If you are using a Microsoft Windows machine, use `set` instead of `export`.

```
export AWS_ACCESS_KEY_ID      = "access_key_from_assume_role"
export AWS_SECRET_ACCESS_KEY  = "secret_key_from_assume_role"
export AWS_SESSION_TOKEN     = "session_token_from_assume_role"
```

Running these commands sets the role session ID of the user visiting your website to `embedding_quicksight_dashboard_role/john.doe@example.com`. The role session ID is made up of the role name from `role-arn` and the `role-session-name` value. Using the unique role session ID for each user ensures that appropriate permissions are set for each user. It also prevents any throttling of user access. *Throttling* is a security feature that prevents the same user from accessing Amazon Quick Sight from multiple locations.

The role session ID also becomes the user name in Amazon Quick Sight. You can use this pattern to provision your users in Amazon Quick Sight ahead of time, or to provision them the first time they access the dashboard.

The following example shows the CLI command that you can use to provision a user. For more information about [RegisterUser](#), [DescribeUser](#), and other Amazon Quick Sight API operations, see the [Amazon Quick Sight API reference](#).

```
aws quicksight register-user \
  --aws-account-id 111122223333 \
  --namespace default \
  --identity-type IAM \
  --iam-arn "arn:aws:iam::111122223333:role/embedding_quicksight_dashboard_role" \
```

```
--user-role READER \  
--user-name jhnd \  
--session-name "john.doe@example.com" \  
--email john.doe@example.com \  
--region us-east-1 \  
--custom-permissions-name TeamA1
```

If the user is authenticated through Microsoft AD, you don't need to use `RegisterUser` to set them up. Instead, they should be automatically subscribed the first time they access Amazon Quick Sight. For Microsoft AD users, you can use `DescribeUser` to get the user ARN.

The first time a user accesses Amazon Quick Sight, you can also add this user to the group that the dashboard is shared with. The following example shows the CLI command to add a user to a group.

```
aws quicksight create-group-membership \  
  --aws-account-id=111122223333 \  
  --namespace=default \  
  --group-name=financeusers \  
  --member-name="embedding_quicksight_dashboard_role/john.doe@example.com"
```

You now have a user of your app who is also a user of Amazon Quick Sight, and who has access to the dashboard.

Finally, to get a signed URL for the dashboard, call `get-dashboard-embed-url` from the app server. This returns the embeddable dashboard URL. The following example shows how to get the URL for an embedded dashboard using a server-side call for users authenticated through AWS Managed Microsoft AD or IAM Identity Center.

```
aws quicksight get-dashboard-embed-url \  
  --aws-account-id 111122223333 \  
  --dashboard-id 1a1ac2b2-3fc3-4b44-5e5d-c6db6778df89 \  
  --identity-type IAM \  
  --session-lifetime-in-minutes 30 \  
  --undo-redo-disabled true \  
  --reset-disabled true \  
  --state-persistence-enabled true \  
  --user-arn arn:aws:quicksight:us-east-1:111122223333:user/default/  
embedding_quicksight_dashboard_role/embeddingsession
```


For more information on using this operation, see [GetDashboardEmbedUrl](#). You can use this and other API operations in your own code.

Step 3: Embed the dashboard URL

Important

Amazon Quick Sight has new APIs for embedding analytics: `GenerateEmbedUrlForAnonymousUser` and `GenerateEmbedUrlForRegisteredUser`. You can still use the `GetDashboardEmbedUrl` and `GetSessionEmbedUrl` APIs to embed dashboards and the Amazon Quick Sight console, but they do not contain the latest embedding capabilities. For the latest up-to-date embedding experience, see [Embedding Amazon Quick Sight analytics into your applications](#).

In the following section, you can find out how you can use the [Amazon Quick Sight embedding SDK](#) (JavaScript) to embed the dashboard URL from step 3 in your website or application page. With the SDK, you can do the following:

- Place the dashboard on an HTML page.
- Pass parameters into the dashboard.
- Handle error states with messages that are customized to your application.

Call the `GetDashboardEmbedUrl` API operation to get the URL that you can embed in your app. This URL is valid for 5 minutes, and the resulting session is valid for 10 hours. The API operation provides the URL with an `auth_code` that enables a single-sign on session.

The following shows an example response from `get-dashboard-embed-url`.

```
//The URL returned is over 900 characters. For this example, we've shortened the string
for
//readability and added ellipsis to indicate that it's incomplete.
{
  "Status": "200",
  "EmbedUrl": "https://dashboards.example.com/
embed/620bef10822743fab329fb3751187d2d...",
  "RequestId": "7bee030e-f191-45c4-97fe-d9faf0e03713"
}
```

Embed this dashboard in your webpage by using the [Amazon Quick Sight embedding SDK](#) or by adding this URL into an iframe. If you set a fixed height and width number (in pixels), Amazon Quick Sight uses those and doesn't change your visual as your window resizes. If you set a relative percent height and width, Amazon Quick Sight provides a responsive layout that is modified as your window size changes. By using the Amazon Quick Sight Embedding SDK, you can also control parameters within the dashboard and receive callbacks in terms of page load completion and errors.

The following example shows how to use the generated URL. This code is generated on your app server.

```
<!DOCTYPE html>
<html>

<head>
  <title>Basic Embed</title>

  <script src="./quicksight-embedding-js-sdk.min.js"></script>
  <script type="text/javascript">
    var dashboard;

    function embedDashboard() {
      var containerDiv = document.getElementById("embeddingContainer");
      var options = {
        // replace this dummy url with the one generated via embedding API
        url: "https://us-east-1.quicksight.aws.amazon.com/sn/dashboards/
dashboardId?isauthcode=true&identityprovider=quicksight&code=authcode",
        container: containerDiv,
        scrolling: "no",
        height: "700px",
        width: "1000px",
        footerPaddingEnabled: true
      };
      dashboard = QuickSightEmbedding.embedDashboard(options);
    }
  </script>
</head>

<body onload="embedDashboard()">
  <div id="embeddingContainer"></div>
</body>
```

```
</html>
```

For this example to work, make sure to use the Amazon Quick Sight Embedding SDK to load the embedded dashboard on your website using JavaScript. To get your copy, do one of the following:

- Download the [Amazon Quick Sight embedding SDK](#) from GitHub. This repository is maintained by a group of Amazon Quick Sight developers.
- Download the latest embedding SDK version from <https://www.npmjs.com/package/amazon-quicksight-embedding-sdk>.
- If you use npm for JavaScript dependencies, download and install it by running the following command.

```
npm install amazon-quicksight-embedding-sdk
```

Embedding the Amazon Quick Sight console using GetSessionEmbedUrl (old API)

Important

Amazon Quick Sight has new APIs for embedding analytics:

`GenerateEmbedUrlForAnonymousUser` and `GenerateEmbedUrlForRegisteredUser`.

You can still use the `GetDashboardEmbedUrl` and `GetSessionEmbedUrl` APIs to embed dashboards and the Amazon Quick Sight console, but they do not contain the latest embedding capabilities. For the latest up-to-date embedding experience, see [Embedding Amazon Quick Sight analytics into your applications](#).

Applies to: Enterprise Edition

Intended audience: Amazon Quick developers

In the following sections, you can find detailed information on how to provide the Amazon Quick Sight console experience in a custom-branded authoring portal for registered users using the `GetSessionEmbedUrl` API.

Topics

- [Step 1: Set up permissions](#)
- [Step 2: Get the URL with the authentication code attached](#)
- [Step 3: Embed the console session URL](#)

Step 1: Set up permissions

Important

Amazon Quick Sight has new APIs for embedding analytics:

`GenerateEmbedUrlForAnonymousUser` and `GenerateEmbedUrlForRegisteredUser`. You can still use the `GetDashboardEmbedUrl` and `GetSessionEmbedUrl` APIs to embed dashboards and the Amazon Quick Sight console, but they do not contain the latest embedding capabilities. For the latest up-to-date embedding experience, see [Embedding Amazon Quick Sight analytics into your applications](#).

In the following section, you can find out how to set up permissions for the backend application or web server. This task requires administrative access to IAM.

Each user who accesses a Amazon Quick Sight assumes a role that gives them Amazon Quick Sight access and permissions to the console session. To make this possible, create an IAM role in your AWS account. Associate an IAM policy with the role to provide permissions to any user who assumes it. Add `quicksight:RegisterUser` permissions to ensure that the reader can access Amazon Quick Sight in a read-only fashion, and not have access to any other data or creation capability. The IAM role also needs to provide permissions to retrieve console session URLs. For this, you add `quicksight:GetSessionEmbedUrl`.

The following sample policy provides these permissions for use with `IdentityType=IAM`.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": "quicksight:RegisterUser",
```

```
    "Resource": "*",
    "Effect": "Allow"
  },
  {
    "Action": "quicksight:GetSessionEmbedUrl",
    "Resource": "*",
    "Effect": "Allow"
  }
]
```

The following sample policy provides permission to retrieve a console session URL. You use the policy without `quicksight:RegisterUser` if you are creating users before they access an embedded session.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "quicksight:GetSessionEmbedUrl"
      ],
      "Resource": "*"
    }
  ]
}
```

If you use `QUICKSIGHT` as your `identityType` and provide the user's Amazon Resource Name (ARN), you also need to allow the `quicksight:GetAuthCode` action in your policy. The following sample policy provides this permission.

JSON

```
{
  "Version": "2012-10-17",
```

```
"Statement": [
  {
    "Effect": "Allow",
    "Action": [
      "quicksight:GetSessionEmbedUrl",
      "quicksight:GetAuthCode"
    ],
    "Resource": "*"
  }
]
```

Your application's IAM identity must have a trust policy associated with it to allow access to the role that you just created. This means that when a user accesses your application, your application can assume the role on the user's behalf and provision the user in Amazon Quick Sight. The following example shows a role called `embedding_quicksight_console_session_role`, which has the sample policy preceding as its resource.

For more information regarding trust policies for OpenID Connect or SAML authentication, see the following sections of the *IAM User Guide*:

- [Creating a role for web identity or OpenID Connect federation \(console\)](#)
- [Creating a role for SAML 2.0 federation \(console\)](#)

Step 2: Get the URL with the authentication code attached

Important

Amazon Quick Sight has new APIs for embedding analytics: `GenerateEmbedUrlForAnonymousUser` and `GenerateEmbedUrlForRegisteredUser`. You can still use the `GetDashboardEmbedUrl` and `GetSessionEmbedUrl` APIs to embed dashboards and the Amazon Quick Sight console, but they do not contain the latest embedding capabilities. For the latest up-to-date embedding experience, see [Embedding Amazon Quick Sight analytics into your applications](#).

In the following section, you can find out how to authenticate your user and get the embeddable console session URL on your application server.

When a user accesses your app, the app assumes the IAM role on the user's behalf. Then it adds the user to Amazon Quick Sight, if that user doesn't already exist. Next, it passes an identifier as the unique role session ID.

Performing the described steps ensures that each viewer of the console session is uniquely provisioned in Amazon Quick Sight. It also enforces per-user settings, such as the row-level security and dynamic defaults for parameters.

The following examples perform the IAM authentication on the user's behalf. This code runs on your app server.

Java

```
import com.amazonaws.auth.AWSCredentials;
import com.amazonaws.auth.BasicAWSCredentials;
import com.amazonaws.auth.AWSCredentialsProvider;
import com.amazonaws.regions.Regions;
import com.amazonaws.services.quicksight.AmazonQuickSight;
import com.amazonaws.services.quicksight.AmazonQuickSightClientBuilder;
import com.amazonaws.services.quicksight.model.GetSessionEmbedUrlRequest;
import com.amazonaws.services.quicksight.model.GetSessionEmbedUrlResult;

/**
 * Class to call QuickSight AWS SDK to get url for session embedding.
 */
public class GetSessionEmbedUrlQSAuth {

    private final AmazonQuickSight quickSightClient;

    public GetSessionEmbedUrlQSAuth() {
        this.quickSightClient = AmazonQuickSightClientBuilder
            .standard()
            .withRegion(Regions.US_EAST_1.getName())
            .withCredentials(new AWSCredentialsProvider() {
                @Override
                public AWSCredentials getCredentials() {
                    // provide actual IAM access key and secret
                    key here
                    return new BasicAWSCredentials("access-
key", "secret-key");
                }
            })
            .build();
    }

    @Override
```

```

        public void refresh() {}
    }

    )
    .build();
}

public String getQuicksightEmbedUrl(
    final String accountId, // YOUR AWS ACCOUNT ID
    final String userArn // REGISTERED USER ARN TO USE FOR EMBEDDING.
    REFER TO GETEMBEDURL SECTION IN DEV PORTAL TO FIND OUT HOW TO GET USER ARN FOR A
    QUICKSIGHT USER
    ) throws Exception {
    GetSessionEmbedUrlRequest getSessionEmbedUrlRequest = new
    GetSessionEmbedUrlRequest()
        .withAwsAccountId(accountId)
        .withEntryPoint("/start")
        .withUserArn(userArn);

    GetSessionEmbedUrlResult sessionEmbedUrl =
    quickSightClient.getSessionEmbedUrl(getSessionEmbedUrlRequest);

    return sessionEmbedUrl.getEmbedUrl();
}
}

```

JavaScript

```

global.fetch = require('node-fetch');
const AWS = require('aws-sdk');

function getSessionEmbedURL(
    accountId, // YOUR AWS ACCOUNT ID
    userArn, // REGISTERED USER ARN TO USE FOR EMBEDDING. REFER TO GETEMBEDURL
    SECTION IN DEV PORTAL TO FIND OUT HOW TO GET USER ARN FOR A QUICKSIGHT USER
    getEmbedUrlCallback, // GETEMBEDURL SUCCESS CALLBACK METHOD
    errorCallback // GETEMBEDURL ERROR CALLBACK METHOD
    ) {
    const getSessionParams = {
        AwsAccountId: accountId,
        EntryPoint: "/start",
        UserArn: userArn,
        SessionLifetimeInMinutes: 600,
    };
};

```



```

const quicksightGetSession = new AWS.QuickSight({
  region: process.env.AWS_REGION,
});

quicksightGetSession.getSessionEmbedUrl(getSessionParams, function(err, data) {
  if (err) {
    console.log(err, err.stack);
    errorCallback(err);
  } else {
    const result = {
      "statusCode": 200,
      "headers": {
        "Access-Control-Allow-Origin": "*", // USE YOUR WEBSITE DOMAIN
        "Access-Control-Allow-Headers": "Content-Type"
      },
      "body": JSON.stringify(data),
      "isBase64Encoded": false
    }
    getEmbedUrlCallback(result);
  }
});
}

```

Python3

```

import json
import boto3
from botocore.exceptions import ClientError
import time

# Create QuickSight and STS clients
qs = boto3.client('quicksight', region_name='us-east-1')
sts = boto3.client('sts')

# Function to generate embedded URL
# accountId: YOUR AWS ACCOUNT ID
# userArn: REGISTERED USER ARN TO USE FOR EMBEDDING. REFER TO GETEMBEDURL SECTION IN
# DEV PORTAL TO FIND OUT HOW TO GET USER ARN FOR A QUICKSIGHT USER
def getSessionEmbedURL(accountId, userArn):
    try:
        response = qs.get_session_embed_url(

```

```

        AwsAccountId = accountId,
        EntryPoint = "/start",
        UserArn = userArn,
        SessionLifetimeInMinutes = 600
    )

    return {
        'statusCode': 200,
        'headers': {"Access-Control-Allow-Origin": "*", "Access-Control-Allow-Headers": "Content-Type"},
        'body': json.dumps(response),
        'isBase64Encoded': bool('false')
    }
except ClientError as e:
    print(e)
    return "Error generating embeddedURL: " + str(e)

```

Node.js

The following example shows the JavaScript (Node.js) that you can use on the app server to get the URL for the embedded console session. You can use this URL in your website or app to display the console session.

Example

```

const AWS = require('aws-sdk');
const https = require('https');

var quicksight = new AWS.Service({
    apiConfig: require('./quicksight-2018-04-01.min.json'),
    region: 'us-east-1',
});

quicksight.getSessionEmbedUrl({
    'AwsAccountId': '111122223333',
    'EntryPoint': 'https://url-for-console-page-to-open',
    'SessionLifetimeInMinutes': 600,
    'UserArn': 'USER_ARN'
}, function(err, data) {
    console.log('Errors: ');
    console.log(err);
    console.log('Response: ');

```

```
        console.log(data);
    });
```

Example

//The URL returned is over 900 characters. For this example, we've shortened the string for

```
    //readability and added ellipsis to indicate that it's incomplete.
    { Status: 200,
      EmbedUrl: 'https://dashboards.example.com/
embed/620bef10822743fab329fb3751187d2d...
      RequestId: '7bee030e-f191-45c4-97fe-d9faf0e03713' }
```

.NET/C#

The following example shows the .NET/C# code that you can use on the app server to get the URL for the embedded console session. You can use this URL in your website or app to display the console.

Example

```
var client = new AmazonQuickSightClient(
    AccessKey,
    SecretAccessKey,
    sessionToken,
    Amazon.RegionEndpoint.USEast1);
try
{
    Console.WriteLine(
        client.GetSessionEmbedUrlAsync(new GetSessionEmbedUrlRequest
        {
            'AwsAccountId': '111122223333',
            'EntryPoint': 'https://url-for-console-page-to-open',
            'SessionLifetimeInMinutes': 600,
            'UserArn': 'USER_ARN'
            AwsAccountId = 111122223333,
            EntryPoint = https://url-for-console-page-to-open,
            SessionLifetimeInMinutes = 600,
            UserArn = 'USER_ARN'
        }).Result.EmbedUrl
    );
```

```
    } catch (Exception ex) {  
        Console.WriteLine(ex.Message);  
    }
```

AWS CLI

To assume the role, choose one of the following AWS Security Token Service (AWS STS) API operations:

- [AssumeRole](#) – Use this operation when you are using an IAM identity to assume the role.
- [AssumeRoleWithWebIdentity](#) – Use this operation when you are using a web identity provider to authenticate your user.
- [AssumeRoleWithSaml](#) – Use this operation when you are using SAML to authenticate your users.

The following example shows the CLI command to set the IAM role. The role needs to have permissions enabled for `quicksight:GetSessionEmbedUrl`. If you are taking a just-in-time approach to add users when they first open Amazon Quick Sight, the role also needs permissions enabled for `quicksight:RegisterUser`.

```
aws sts assume-role \  
    --role-arn "arn:aws:iam::111122223333:role/embedding_quicksight_dashboard_role"  
 \  
    --role-session-name john.doe@example.com
```

The `assume-role` operation returns three output parameters: the access key, the secret key, and the session token.

Note

If you get an `ExpiredToken` error when calling the `AssumeRole` operation, this is probably because the previous `SESSION_TOKEN` is still in the environment variables. Clear this by setting the following variables:

- `AWS_ACCESS_KEY_ID`
- `AWS_SECRET_ACCESS_KEY`
- `AWS_SESSION_TOKEN`

The following example shows how to set these three parameters in the CLI. If you are using a Microsoft Windows machine, use `set` instead of `export`.

```
export AWS_ACCESS_KEY_ID      = "access_key_from_assume_role"
export AWS_SECRET_ACCESS_KEY  = "secret_key_from_assume_role"
export AWS_SESSION_TOKEN      = "session_token_from_assume_role"
```

Running these commands sets the role session ID of the user visiting your website to `embedding_quicksight_console_session_role/john.doe@example.com`. The role session ID is made up of the role name from `role-arn` and the `role-session-name` value. Using the unique role session ID for each user ensures that appropriate permissions are set for each user. It also prevents any throttling of user access. Throttling is a security feature that prevents the same user from accessing Amazon Quick Sight from multiple locations.

The role session ID also becomes the user name in Amazon Quick Sight. You can use this pattern to provision your users in Amazon Quick Sight ahead of time, or to provision them the first time they access a console session.

The following example shows the CLI command that you can use to provision a user. For more information about [RegisterUser](#), [DescribeUser](#), and other Amazon Quick Sight API operations, see the [Amazon Quick Sight API reference](#).

```
aws quicksight register-user \
  --aws-account-id 111122223333 \
  --namespace default \
  --identity-type IAM \
  --iam-arn "arn:aws:iam::111122223333:role/embedding_quicksight_dashboard_role" \
  --user-role READER \
  --user-name jhnd \
  --session-name "john.doe@example.com" \
  --email john.doe@example.com \
  --region us-east-1 \
  --custom-permissions-name TeamA1
```

If the user is authenticated through Microsoft AD, you don't need to use `RegisterUser` to set them up. Instead, they should be automatically subscribed the first time they access Amazon Quick Sight. For Microsoft AD users, you can use `DescribeUser` to get the user ARN.

The first time a user accesses Amazon Quick Sight, you can also add this user to the appropriate group. The following example shows the CLI command to add a user to a group.

```
aws quicksight create-group-membership \  
  --aws-account-id=111122223333 \  
  --namespace=default \  
  --group-name=financeusers \  
  --member-name="embedding_quicksight_dashboard_role/john.doe@example.com"
```

You now have a user of your app who is also a user of Amazon Quick Sight, and who has access to the Amazon Quick Sight console session.

Finally, to get a signed URL for the console session, call `get-session-embed-url` from the app server. This returns the embeddable console session URL. The following example shows how to get the URL for an embedded console session using a server-side call for users authenticated through AWS Managed Microsoft AD or Single Sign-on (IAM Identity Center).

```
aws quicksight get-dashboard-embed-url \  
  --aws-account-id 111122223333 \  
  --entry-point the-url-for--the-console-session \  
  --session-lifetime-in-minutes 600 \  
  --user-arn arn:aws:quicksight:us-east-1:111122223333:user/  
default/embedding_quicksight_dashboard_role/embeddingession
```

For more information on using this operation, see [GetSessionEmbedUrl](#). You can use this and other API operations in your own code.

Step 3: Embed the console session URL

Important

Amazon Quick Sight has new APIs for embedding analytics: `GenerateEmbedUrlForAnonymousUser` and `GenerateEmbedUrlForRegisteredUser`. You can still use the `GetDashboardEmbedUrl` and `GetSessionEmbedUrl` APIs to embed dashboards and the Amazon Quick Sight console, but they do not contain the latest embedding capabilities. For the latest up-to-date embedding experience, see [Embedding Amazon Quick Sight analytics into your applications](#).

In the following section, you can find out how you can use the [Amazon Quick Sight embedding SDK](#) (JavaScript) to embed the console session URL from step 3 in your website or application page. With the SDK, you can do the following:

- Place the console session on an HTML page.
- Pass parameters into the console session.
- Handle error states with messages that are customized to your application.

Call the `GetSessionEmbedUrl` API operation to get the URL that you can embed in your app. This URL is valid for 5 minutes, and the resulting session is valid for 10 hours. The API operation provides the URL with an `auth_code` that enables a single-sign on session.

The following shows an example response from `get-dashboard-embed-url`.

```
//The URL returned is over 900 characters. For this example, we've shortened the string
for
//readability and added ellipsis to indicate that it's incomplete.
{
  "Status": "200",
  "EmbedUrl": "https: //dashboards.example.com/
embed/620bef10822743fab329fb3751187d2d...",
  "RequestId": "7bee030e-f191-45c4-97fe-d9faf0e03713"
}
```

Embed this console session in your webpage by using the Amazon Quick Sight [Embedding SDK](#) or by adding this URL into an `iframe`. If you set a fixed height and width number (in pixels), Amazon Quick Sight uses those and doesn't change your visual as your window resizes. If you set a relative percent height and width, Amazon Quick Sight provides a responsive layout that is modified as your window size changes. By using the Amazon Quick Sight Embedding SDK, you can also control parameters within the console session and receive callbacks in terms of page load completion and errors.

The following example shows how to use the generated URL. This code is generated on your app server.

```
<!DOCTYPE html>
<html>

<head>
  <title>Basic Embed</title>

  <script src="./quicksight-embedding-js-sdk.min.js"></script>
  <script type="text/javascript">
    var dashboard;
```

```
function embedDashboard() {
    var containerDiv = document.getElementById("embeddingContainer");
    var options = {
        // replace this dummy url with the one generated via embedding API
        url: "https://us-east-1.quicksight.aws.amazon.com/sn/dashboards/
dashboardId?isauthcode=true&identityprovider=quicksight&code=authcode",
        container: containerDiv,
        scrolling: "no",
        height: "700px",
        width: "1000px",
        footerPaddingEnabled: true
    };
    dashboard = QuickSightEmbedding.embedDashboard(options);
}
</script>
</head>

<body onload="embedDashboard()">
    <div id="embeddingContainer"></div>
</body>

</html>
```

For this example to work, make sure to use the Amazon Quick Sight Embedding SDK to load the embedded console session on your website using JavaScript. To get your copy, do one of the following:

- Download the [Amazon Quick Sight embedding SDK](#) from GitHub. This repository is maintained by a group of Amazon Quick Sight developers.
- Download the latest embedding SDK version from <https://www.npmjs.com/package/amazon-quicksight-embedding-sdk>.
- If you use npm for JavaScript dependencies, download and install it by running the following command.

```
npm install amazon-quicksight-embedding-sdk
```


Using Amazon Quick Research

Amazon Quick Research is a feature of Amazon Quick that enables you to conduct comprehensive research by analyzing multiple data sources and generating detailed reports. Quick Research uses artificial intelligence to help you gather, analyze, and synthesize information from various sources including web search, uploaded files, connected data spaces, knowledge bases, actions, and third-party data providers.

With Quick Research, you can define research objectives, select relevant data sources, and receive AI-generated research reports with proper citations and source tracing. This helps you make informed decisions based on comprehensive analysis of available information.

What you can do with Quick Research

Quick Research supports a wide range of research activities and use cases:

- Market research and competitive analysis
- Literature reviews and academic research
- Business intelligence and trend analysis
- Product research and feature comparison
- Policy and regulatory research
- Technical documentation and best practices research
- Customer insights and feedback analysis
- Industry reports and white paper synthesis

Quick Research helps you save time by automatically gathering information from multiple sources, analyzing the content, and presenting findings in a structured format with proper citations and source references.

Starting new research

Creating a new research project in Amazon Quick Research involves a structured workflow that guides you from initial objective setting through final report generation. This process ensures comprehensive research coverage while maintaining focus on your specific goals.

Enter research objective

The first step in creating a new research project is to define your research objective. A well-written research objective helps Quick Research understand what information you're looking for and guides the AI in selecting relevant sources and generating focused results.

To enter your research objective

1. In the main navigation bar, choose **Research**.
2. In the Quick Research interface, choose **New Research**.
3. In the **Research objective** field, enter a clear description of what you want to research.

Make your research goal specific and detailed. For example, include the following:

- What you are researching and why
- What industry, region, and timeframe
- Any data sources like dashboards, documents, or internet resources
- Key focus areas such as metrics, competitors, and trends

Share your research objective in a few sentences to describe your goal, frameworks, desired outcomes and any other essential information including your industry context and any specific requirements or constraints.

Example: "Evaluate renewable energy investment opportunities for our corporate sustainability portfolio to support Q1 2025 capital allocation decisions. We're allocating \$50M across North American markets with a 7-year investment horizon. Focus on solar, wind, and battery storage sectors from 2023-2024 performance data. Key metrics: IRR projections, regulatory incentive impacts, market growth trends, and risk-adjusted returns. Cost sensitivity is critical - prioritize solutions under \$2M per MW capacity."

You can also reference internal data sources in your objective. For example:

"Primarily build research on our internal investment analysis from Document: 'Renewable_Portfolio_Performance_2024.pdf' in Knowledge base: 'Investment Research', which is added in Space: 'Financial Analysis', and Chart: 'Investment Performance Trends' in Dashboard: 'Sustainability Portfolio Dashboard'. Use web research to supplement with recent policy changes, emerging technology costs, and competitor strategies not covered in our internal data."

4. (Optional) Choose the **Research mode** button to select a research mode. The following modes are available:
 - **Fast** – Completes research in approximately 4-7 minutes. This mode is suitable for quick insights and time-sensitive research needs.
 - **Deep** – Completes research in approximately 20-30 minutes. This mode provides more thorough analysis and is suitable for complex research topics that require deeper investigation.
5. After entering your research objective, choose what materials to use in **Research materials**, then choose **Review plan** to review a draft plan, or choose **Start researching** to begin research immediately.

Tips for writing a good research objective

Follow these guidelines to create effective research objectives:

- Be specific about what you want to learn or understand
- Include relevant context about your industry, market, or domain
- Specify the time frame if relevant (for example, "recent trends," "2024 data")
- Mention the intended use of the research (for example, "for strategic planning," "to inform product decisions")
- Use clear, concise language and avoid ambiguous terms
- Include any specific aspects or angles you want to explore

Example of a well-written research objective: "Analyze the current state of artificial intelligence adoption in healthcare, focusing on diagnostic imaging applications, regulatory challenges, and market opportunities for 2024-2025."

Select research materials

After defining your research objective, you can select the data sources that Quick Research will use to gather information. Set your preferences and control what sources Quick Research uses. Include files, web search preferences, company assets, and premium and public datasets.

To select research materials

1. In the material selection interface, review the available data source options.

2. Select the types of materials you want to include:

- a. Toggle **Web search** to include online sources
- b. Choose **File upload** to add specific documents

For best results, include key documents that are most relevant to your research goals. Adding more materials may increase research time and lead to less focused insights.

If you're working with larger collections, consider using Spaces to organize and optimize extensive research materials for better results.

- c. Choose **Quick assets** to include data spaces, dashboards, knowledge bases, and integrations in your research. For more information about creating and managing data spaces, see [Organize and collaborate using Amazon Quick spaces](#).
 - d. Choose **Third party data** to incorporate data from premium datasets and trusted public data sources to enhance your research. For more information, see [the section called "Using third party data for research"](#).
3. Choose **Review plan** to review a draft research plan before starting, or choose **Start researching** to begin research immediately.

Web search

Enable web search to allow Quick Research to gather information from publicly available online sources. This includes academic papers, industry reports, news articles, and other relevant web content related to your research objective.

Amazon Quick uses the internet to enhance your results. Web search queries may be processed in a different AWS Region than your primary Region. For more information, see [Cross-Region calls for web search](#).

To prioritize or avoid certain websites, you can expand the web search section to enter a list of preferred websites and a list of websites to avoid. Both fields are optional and can take a maximum of 3,500 characters.

We recommend providing a list of web domains (such as `example.com`), but you can also include types of websites such as `government websites` or `blogs`. If you enter a website address like `example.com/path/to/specific/page`, it will be shortened to `example.com`, so you don't need to enter multiple websites for a single domain.

Note

Adding a domain to the list of preferred websites doesn't guarantee that the website will be used in research. A website might not be used include: the site is behind a paywall, the site doesn't allow agents to access it, or the site's content is found to be less relevant than other sites.

Uploaded files

Upload specific documents, PDFs, spreadsheets, or other files that you want Quick Research to analyze as part of your research. This is useful when you have particular sources or documents that are directly relevant to your research objective.

To upload files

1. Choose **Upload** to open the **File Uploads** window.
2. Add files using one of the following methods:
 - a. Choose **Add Files** to navigate to your files in File Explorer
 - b. Drag and drop up to 20 files into the upload area

If you need to include more than 20 files in your research, consider organizing them in Spaces instead.

3. When you're finished selecting files, choose **Done**.

Accepted file formats include .pdf, .txt, .doc, .xls, and .csv files. The file size limit is 25MB per file.

Quick assets

Connect Quick Research to your existing data spaces to include internal documents, reports, knowledge bases, and integrations in your research. You can combine external web sources with your organization's proprietary information. Quick Research can also fetch data through Quick connectors, enabling it to pull live data from connected systems during the research process.

For optimal results, consider selecting the most relevant assets for your specific research question. Choose **Browse** to add assets from Quick.

To add Quick assets

1. Choose **Browse** to open the **Add assets** window in the right pane.
2. Select from the available resource type tabs:
 - a. **Recent** - Recently accessed assets
 - b. **Space** - Available data spaces (maximum of 2 selections)
 - c. **Dashboard** - Dashboard resources (maximum of 2 selections)
 - d. **Topic** - Topic resources (maximum of 2 selections)
 - e. **Knowledgebase** - Knowledge base resources
 - f. **Integration** - Connectors that can fetch live data from connected systems
3. Review the available assets, which display the Name, Owner, and Last Modified date and time for each asset.
4. Select the assets you want to include in your research.
5. When you're finished specifying all assets, choose **Add**.

Review and finalize research plan

Before Quick Research begins gathering and analyzing information, you have the opportunity to review and refine your research plan. This step helps ensure that the research will meet your specific needs and expectations.

To review and finalize your research plan

1. Review the AI-generated suggestions in the **Revise Plan** pane on the right for improving your research plan.
2. (Optional) To make revisions, enter your changes in the **How should I change the plan?** field. You can add comments and notes to provide additional context or specific instructions for the AI, such as clarifying particular aspects you want emphasized, specifying exclusions, or providing domain-specific guidance that will help generate more relevant results.
3. Choose **Revise plan** to update the plan with your revisions.
4. Choose **Start researching** to begin the analysis process.

Research suggestions

Quick Research analyzes your research objective and selected materials to provide intelligent suggestions for improving your research plan. These suggestions appear in the **Revise Plan** pane and might include additional data sources, refined search terms, or alternative approaches to your research question.

For example, if your research objective is about sustainable packaging for the food industry, suggestions might include specifying geographic regions for regulatory analysis, identifying specific food types requiring temperature maintenance, or providing budget constraints for cost comparison analysis.

Revision history

The revision history section tracks all changes made to your research plan during the review process. This allows you to see what modifications have been applied and provides a record of how your research plan has evolved.

View research report

Once you select to create the research, Quick Research will work through a series of research tasks. Research typically takes 20-40 minutes to complete. The window displays the percentage complete along with the specific tasks Quick Research is working on and which sources are being used for each research task.

After Quick Research completes its analysis, you can view the comprehensive research report that synthesizes information from all selected sources. The report presents findings in a structured format with supporting evidence and source references.

You can use the **Download** button at the top to download the report as a PDF or Word document. There is also a **Summarize** button in the upper right. For more information, see [the section called "Summarizing the research"](#). You can also select **Reading mode** to hide sidebars and focus on the report content.

When you're ready to share your research with others, use the **Share** button. For more information, see [the section called "Sharing the research"](#). If you want to refine the research further before sharing, you can make updates and iterate on the report. For more information, see [the section called "Updating research"](#).

In the left pane, under your research objective you can see a **Topics** tab which lists and allows you to navigate through the major sections of the report.

The **Research** pane on the right of the window accepts comments up to 400 characters. You can now add comments that will be queued for revision, as they require deeper investigation by Quick Research.

Viewing citations

Throughout the report, numbered citations are present. Clicking on a citation shows you a popup window with the source article and a hyperlink to the source page.

Understanding statements

You can also click the **Understand the statement** icon (three horizontal lines and a plus) to launch an explanation window that shows how a statement in the report was determined, including a summary of the evidence and a breakdown of the conclusion.

Updating research

Research is an iterative process that often requires refinement and updates based on new findings or changing requirements. Quick Research provides flexible tools to modify your research objectives, incorporate additional sources, and track changes through version management.

Making updates

After reviewing your initial research report, you can make updates to refine and improve the results. Quick Research provides multiple ways to provide feedback and request modifications to better meet your needs.

Once you've added all your comments to the queue in the Research pane, the **Revise** button becomes available in the upper right. When you choose **Revise**, you see a message that "Review revision started". Quick Research goes back to work through its research tasks with your new comments in mind, analyzing the existing report's structure and content to determine how to apply your comments.

The version increments to Version 2 when the revision is complete.

In-line feedback

Provide specific feedback on individual sections or findings within your research report by highlighting the word or section you want to comment on. In-line feedback allows you to request clarification, additional detail, or corrections for specific parts of the report without affecting the entire document.

High level feedback

Submit broader feedback about the overall direction, focus, or approach of your research. High-level feedback is useful when you want to adjust the research scope, change the emphasis, or request additional analysis across the entire report.

Research versions

Quick Research maintains versions of your research reports as you make updates and refinements. This version control system allows you to track changes, compare different iterations, and revert to previous versions if needed.

Each time you submit feedback or request updates, Quick Research creates a new version of your report while preserving the previous versions. You can view the version history, compare changes between versions, and choose which version to use as your final report.

Version management helps you maintain a clear audit trail of your research process and ensures that you don't lose valuable insights from earlier iterations of your work.

Summarizing the research

After completing your research report, you can create focused summaries for different audiences. To access the summarizing feature, choose the **Summarize** button in the upper right of your research report.

Quick Research provides multiple summary formats to help you present your findings in different contexts and for various audiences. Choose the summary type that best fits your intended use and audience needs.

In the **Options** tab, you can choose to create a custom summary by entering up to 5,000 characters about how you would like to summarize, or select one of two preset summary styles:

- *Executive summary* - VP-oriented tone, 2-page maximum, no citations

- *General share out* - Business-friendly tone, 6-page maximum, essential citations

Once you choose a summary style, you can choose **Begin** to start producing the summary. The **Previous** tab shows previous summary reports. You can download the summary as PDF or Word using the **Download** button at the top.

Custom summary

Create a tailored summary by entering up to 5,000 characters about how you would like to summarize your research. Custom summaries allow you to specify the key points, length, and focus areas you want emphasized for your particular use case.

Executive summary

Generate a concise executive summary with a VP-oriented tone, limited to 2 pages maximum with no citations. This format presents the most important findings and recommendations suitable for senior leadership and decision-makers.

General share out

Create a comprehensive summary document with a business-friendly tone, up to 6 pages maximum with essential citations included. This format balances depth with accessibility for team presentations and stakeholder briefings.

Sharing the research

After completing your research, you can download it in various formats or share it with others. Amazon Quick Research provides flexible options for both downloading and sharing your research findings.

Download the research

You can export your research in multiple formats to suit different needs and workflows.

Export to PDF

PDF format is ideal for sharing formatted documents that maintain their appearance across different devices and platforms.

To export your research to PDF:

- Navigate to your completed research report
- Choose *Download* from the options menu
- Select *PDF* as the export format
- The PDF file will be generated and downloaded to your device

Export to Word

Word format allows for further editing and collaboration using Microsoft Word or compatible applications.

To export your research to Word format:

- Navigate to your completed research report
- Choose *Download* from the options menu
- Select *Word* as the export format
- The Word document will be generated and downloaded to your device

Sharing the research with stakeholders

Quick Research provides several ways to share your research with colleagues and stakeholders, enabling collaboration and knowledge distribution.

To share your research:

- Open the research report you want to share
- Choose *Share* from the options menu
- Select your preferred sharing method from the available options
- Configure sharing permissions and access levels as needed
- Send the share link or invitation to your intended recipients

Shared research maintains all formatting, citations, and interactive elements, allowing recipients to view the complete research experience.

Automating research with Quick Flows

Amazon Quick Research can be integrated into Quick Flows to automate your research workflows. This allows you to standardize research processes, schedule recurring research reports, and share proven research methods across your organization.

When to use research in Quick Flows

Consider using Quick Research as a flow step when you need to:

- **Standardize research processes** across your team by creating reusable workflows
- **Schedule automated research** that runs at specific times (like every Monday 8am)
- **Trigger downstream actions** based on research findings, such as updating records in Salesforce or creating tasks in Jira

Getting started

To add Quick Research as a step in your flows, you'll configure a research agent that defines your research objective, selects data sources, and optionally accepts user inputs. The research results can then be used to drive subsequent actions in your workflow.

For detailed instructions, see [Research](#).

Using third party data in Amazon Quick Research

Amazon Quick Research integrates with third party data providers, such as FactSet, IDC, and S&P Global. The following table shows the current list of providers. The **Premium** data providers require a license to access their data, while the Public data providers do not require a license.

You can use multiple third party data providers in research. The agent will determine which sources to use based on their relevance.

Data provider	Category	Description
FactSet	Premium	Access FactSet's trusted equity prices, news, events, transcripts, fundamentals and consensus estimates - delivered for AI-driven insights.

Data provider	Category	Description
IDC	Premium	The IDC Connector brings verified technology market intelligence into Quick Research for faster, deeper, and more credible insights.
S&P Global Energy	Premium	S&P Global Energy Data provides AI Ready market insights across commodities and energy, refreshed every 30 minutes with rich metadata.
S&P Global Market Intelligence	Premium	Query S&P Global datasets: Capital IQ Financials, Market Data, Transcripts & securities. Saves time for financial services teams
US Patents	Public	20 years of US patents from the US PTO
PubMed	Public	Millions of biomedical and life sciences articles published in PubMed

Topics

- [Using third party data in Quick Research \(FactSet\)](#)
- [Using third party data in Quick Research \(IDC\)](#)
- [Using third party data in Quick Research \(S&P Global Energy\)](#)
- [Using third party data in Quick Research \(S&P Global Market Intelligence\)](#)
- [Using third party data in Quick Research \(US Patents\)](#)
- [Using third party data in Quick Research \(PubMed\)](#)

Using third party data in Quick Research (FactSet)

Access FactSet's trusted equity prices, fundamentals, events, news, transcripts, pricing and consensus estimates - delivered seamlessly for AI-driven insights. With comprehensive coverage and reliable data, FactSet empowers smarter financial analysis, decision-making, and innovation.

- **Global Prices** - End-of-day pricing data, returns, corporate actions, dividends, and shares outstanding for global equities

- **Fundamentals** - Company financial statements, ratios, and business/geographic segment data from regulatory filings
- **Consensus Estimates** - Earnings estimates, surprise data, and buy/hold/sell ratings consensus
- **Transcripts** – Corporate event data and corresponding transcripts
- **StreetAccount News** – Curated, real-time market-moving news and summaries

Setting up the integration

This integration requires a subscription to FactSet. FactSet will need to authorize your account and someone with an Enterprise subscription to Quick (typically an admin or whomever is responsible for your organization's subscription to FactSet data) will need to set up the integration.

The integration can then be shared with others at your organization who should have access. The instructions below describe these steps in more detail.

To use FactSet in Quick, you (or someone from your organization) will need to:

1. Contact FactSet to get access at AWS.QR.Partnership@factset.com.
2. Create a Quick integration for FactSet.
 - a. Open the Research page and select **New Research**.
 - b. In the **Research Materials** section, find **Third party data** and select **Browse** to display a list of supported third party data integrations.
 - c. Find the integration in the list and select **Connect**. This will display a pop-up with a summary of these setup steps. Select **Continue**.
 - d. Enter a name and description.
 - e. For the client ID, provide the client ID you got from FactSet.
 - f. For the other authentication-related fields you can use the suggestions visible below each input field.
3. Share the integration with anyone else at your organization who should have access.

Using the integration in Quick Research

These instructions assume that you already set up the integration (see details above), or someone at your organization set up the integration and shared it with you.

To use this integration in Quick Research:

1. Navigate to the Quick Research page and select **New Research**.
2. On the left-hand side, under **Research Materials**, there is a section called **Third Party Data**. Select the **Browse** button to display a list of supported integrations.
3. If you need to sign in, there will be a **Sign In** button in the status column. Selecting this button will open a new tab where you can sign in. After you have signed in, navigate back to the tab with Quick Research.
4. Select the check box for the integration you want to use.

Using third party data in Quick Research (IDC)

The IDC Connector integrates IDC's industry-leading technology research and market intelligence directly into Amazon Quick Research, enabling business users to generate deeper, more credible insights in minutes. Through this integration, Quick Research users can securely access IDC's trusted data — including market forecasts, vendor assessments, and emerging technology analyses — alongside their own enterprise data and public web sources.

Setting up the integration

This integration requires a subscription to IDC. IDC will need to enable API access for your account, and someone with an Enterprise subscription to Quick (typically an admin or whomever is responsible for your organization's subscription to IDC) will need to set up the integration and knowledge base. The knowledge base can then be shared with others at your organization who should have access. The instructions below describe these steps in more detail.

Important

Creating an IDC knowledge base incurs storage and indexing costs. You should only create one IDC knowledge base for your organization to avoid unnecessary costs. To avoid users creating duplicate knowledge bases, we recommend not sharing the client ID and client secret with anyone else.

To use IDC in Quick, you (or someone from your organization) will need to:

1. Contact IDC to enable API access at <https://www.idc.com/amazonquickresearch>. IDC will give you a client id and client secret.
2. Create an integration for IDC:

- a. Open the Research page and select **New Research**.
 - b. In the **Research Materials** section, find **Third party data** and select **Browse** to display a list of supported third party data integrations.
 - c. Find the integration in the list and select **Connect**. This will display a pop-up with a summary of these setup steps. Select **Continue**.
 - d. Provide a name and description for the integration.
 - e. Use the client id and client secret you received from IDC.
3. Next, you will be prompted to create a knowledge base. This step could take several hours, depending on how much IDC data needs to be ingested and indexed.
 4. To share the knowledge base with others, scroll to the top of the Integrations page, select the menu icon for IDC, and select **Share**. You can share with individual users or groups. Make sure you only share the knowledge base with users who should have access to IDC data.

Using the integration in Quick Research

These instructions assume that you already set up the integration and knowledge base (see details above), or someone else did this and shared it with you.

To use this integration in Quick Research:

1. Navigate to the Quick Research page and select **New Research**.
2. On the left-hand side, under **Research Materials**, there is a section called **Third Party Data**. Select the **Browse** button to display a list of supported integrations.
3. Select the check box for the integration you want to use.

Using third party data in Quick Research (S&P Global Energy)

S&P Global Energy Data aggregates a vast array of textual content across thousands of documents about commodity and energy markets into an AI Ready data format, updated every 30 minutes. It enables users to ask customized questions on a range of topics such as regulatory challenges, investment opportunities, and trade flows.

Sources include PDF documents, news articles, rationales, commentaries, analyses, and more, all in an LLM-friendly format. This extensive dataset features hundreds of thousands of documents enriched with metadata, embeddings, and citations. The inventory of textual content focuses on different time horizons (from daily news to 1-year outlooks to 20+ year long-term scenario views)

for each of the major industry sectors broken down by subscription packages spanning Oil, Gas, Power, Metals, Clean Energy, Agriculture, Shipping sector, and more.

Users can leverage this integration to answer customized questions, leveraging S&P Global's commodity and energy market expertise. Example prompts include (but not limited to):

- Summarize the current regulatory challenges facing the crude oil industry and how they might impact production and pricing strategies.
- Identify the top investment opportunities in India's wind turbine sector in 2025 and beyond, highlighting emerging technologies or companies that are innovating in this space.
- Discuss recent environmental implications of metals extraction in Africa, and suggest sustainable practices that could be adopted in the industry.
- Summarize the current regulatory challenges facing the US natural gas pipelines in 2025 and how that might impact US gas production through end of this decade.

Setting up the integration

This integration requires a subscription to S&P Global Energy. S&P will need to authorize your account and someone with an Enterprise subscription to Quick (typically an admin or whomever is responsible for your organization's access to S&P data) will need to set up the integration. This person (the 'integration owner') will then share the integration with everyone at your organization who should have access. The instructions below describe these steps in more detail.

To use S&P Global Energy in Quick, you (or someone from your organization) will need to:

1. Contact S&P Global at ai.energy@spglobal.com to ensure your organization subscribes to S&P Global Energy. S&P will provide you a client id and client secret.
 - a. If your organization already subscribes, S&P needs to enable your account so you can access this data in Quick.
 - b. If your organization does not subscribe already, S&P can help you get started with a free trial.
2. Create a Quick integration for S&P Global Energy.
 - a. Open the Research page and select **New Research**.
 - b. In the **Research Materials** section, find **Third party data** and select **Browse** to display a list of supported third party data integrations.
 - c. Find the integration in the list and select **Connect**. This will display a pop-up with a summary of these setup steps. Select **Continue**.

- d. Enter a name and description for the integration. You can enter anything you want to help you identify this integration later.
 - e. Enter the client ID and client secret provided by S&P.
 - f. For the remaining fields, you can use the suggested values displayed below each input field.
3. Share the integration with anyone else at your organization who should have access.

Using the integration in Quick Research

These instructions assume that you already set up the integration (see details above), or someone at your organization set up the integration and shared it with you.

To use this integration in Quick Research:

1. Navigate to the Quick Research page and select **New Research**.
2. On the left-hand side, under **Research Materials**, there is a section called **Third Party Data**. Select the **Browse** button to display a list of supported integrations.
3. If you need to sign in, there will be a **Sign In** button in the status column. Selecting this button will open a new tab where you can sign in. After you have signed in, navigate back to the tab with Quick Research.
4. Select the check box for the integration you want to use.

Using third party data in Quick Research (S&P Global Market Intelligence)

Using S&P Global's MCP Server for the Kensho LLM-Ready API, customers can unlock S&P Global's trusted data and insights using natural language queries, enabling financial professionals to answer complex financial questions and get fast, reliable answers grounded in reliable data from S&P Global.

Query a comprehensive range of S&P Global datasets, including S&P Capital IQ Financials, Market Data, earnings call transcripts, and global securities information. This powerful connector is specifically designed to save time and resources for engineering, product, and business teams in financial services by streamlining access to critical company information, detailed financial statements, historical market data, and comprehensive global securities coverage.

Key capabilities include:

- **Compare Company Financials:** Analyze and benchmark financial performance across companies and time periods. Example: "Compare the gross margins and EBITDA trends for Amazon, Google, and Meta over the past 5 years" to identify competitive positioning and operational efficiency trends.
- **Summarize Earnings Calls:** Extract key insights from earnings transcripts and management commentary. Example: "Summarize the key takeaways from Apple's latest earnings call, focusing on guidance and capital allocation" to quickly understand strategic direction and financial priorities.
- **Deep Financial Research Reports:** Generate comprehensive competitive analysis and market research reports. Example: "Generate a comp table for software companies with \$1B+ revenue and 20%+ growth rates" to identify investment opportunities and market positioning.
- **Develop Due Diligence Briefs:** Develop a due diligence brief in a fraction of the time it would normally take – pulling key financials, generating comps tables, benchmarking peers, and flagging risks in minutes. Example: "Pull key financials for Tesla, generate a peer benchmark analysis, and flag potential risks for an investment committee presentation" to support informed decision-making.

Setting up the integration

This integration requires a subscription to S&P Global Market Intelligence. S&P will need to authorize your account and someone with an Enterprise subscription to Quick (typically an admin or whomever is responsible for your organization's access to this data) will need to set up the integration. The instructions below describe these steps in more detail.

To use S&P Global Market Intelligence in Quick, you (or someone from your organization) will need to:

1. Contact S&P Global to ensure your organization subscribes to S&P Global Market Intelligence data. To learn more about how to access S&P Global data via the Kensho LLM-ready API, email commercial@kensho.com for free trials and pricing, or read more [here](#).
2. Create a Quick integration for S&P Global Market Intelligence.
 - a. Open the Research page and select **New Research**.
 - b. In the **Research Materials** section, find **Third party data** and select **Browse** to display a list of supported third party data integrations.
 - c. Find the integration in the list and select **Connect**. This will display a pop-up with a summary of these setup steps. Select **Continue**.

- d. Provide a name and description for the integration. You can enter anything you like here to help you identify the integration.
3. Share the integration with anyone else at your organization who should have access.

Using the integration in Quick Research

These instructions assume that you already set up the integration (see details above), or someone at your organization set up the integration and shared it with you.

To use this integration in Quick Research:

1. Navigate to the Quick Research page and select **New Research**.
2. On the left-hand side, under **Research Materials**, there is a section called **Third Party Data**. Select the **Browse** button to display a list of supported integrations.
3. If you need to sign in, there will be a **Sign In** button in the status column. Selecting this button will open a new tab where you can sign in. After you have signed in, navigate back to the tab with Quick Research.
4. Select the check box for the integration you want to use.

Using third party data in Quick Research (US Patents)

This integration provides access to 20 years of US patents from the US Patent and Trademark Office (USPTO). For more information, see [USPTO](#).

Setting up the integration

There's no setup required since this is a public dataset.

Using the integration in Quick Research

To use this integration in Quick Research:

1. Navigate to the Quick Research page and select **New Research**.
2. On the left-hand side, under **Research Materials**, there is a section called **Third Party Data**. Select the **Browse** button to display a list of supported integrations.
3. Select the check box for the integration you want to use.

Using third party data in Quick Research (PubMed)

This integration provides access to 10 years of PubMed articles and documents. This includes millions of biomedical and life sciences articles. For more information, see [PubMed](#).

Setting up the integration

There's no setup required since this is a public dataset.

Using the integration in Quick Research

To use this integration in Quick Research:

1. Navigate to the Quick Research page and select **New Research**.
2. On the left-hand side, under **Research Materials**, there is a section called **Third Party Data**. Select the **Browse** button to display a list of supported integrations.
3. Select the check box for the integration you want to use.

Admin controls

Administrators can configure various settings and controls for Quick Research to ensure it meets organizational policies and security requirements.

Turn on/off web search

Administrators can enable or disable web search functionality for their organization. When web search is disabled, Quick Research will only use internal data sources such as uploaded files and connected spaces, ensuring that research remains within the organization's controlled data environment.

This control is useful for organizations with strict data governance policies or when working with sensitive information that should not be supplemented with external web sources.

Status definitions

Quick Research projects progress through various status states during their lifecycle. Understanding these status definitions helps you track the progress of your research and know what actions you can take at each stage.

- *Draft* - Research project is being created and configured but has not yet been submitted for processing
- *In Progress* - Quick Research is actively gathering and analyzing information from your selected sources
- *Review* - Initial research is complete and available for your review and feedback
- *Updating* - Quick Research is processing your feedback and generating an updated version of the report
- *Complete* - Research is finalized and ready for download, sharing, or archiving
- *Error* - An issue occurred during processing that requires attention or retry

You can view the current status of all your research projects from the main Quick Research dashboard.

Using Amazon Quick automations

Quick unifies AI agents for research, business insights, and automation into a single experience, helping business users make better decisions and act on them faster. Quick's automation capabilities handle everything from individual daily tasks to mission-critical enterprise processes, so teams can focus on higher-value work.

Types of automations

Quick offers two automation capabilities designed for different levels of scale and complexity:

- **Amazon Quick Flows** empowers any business user to automate repetitive, routine tasks without needing technical expertise.
- **Amazon Quick Automate** enables power users to orchestrate end-to-end business processes at enterprise scale. Amazon Quick Automate is not available for Free or Plus accounts.

Choosing between Quick Flows and Quick Automate

The right choice depends on the scope and complexity of what you need to automate.

- **User and ownership model.** Choose Quick Flows for personal and team productivity tasks, run by individual users. Choose Quick Automate for centralized, enterprise-wide automations, especially those which need to run independent of an individual user.
- **Workflow scope and duration.** Use Quick Flows for shorter-running tasks such as generating reports, drafting content, or responding to tickets. Use Quick Automate for long-running and high-volume processes such as managing loan applications, processing invoices, and validating new vendor compliance.
- **Customization.** Quick Flows offers a simple but powerful set of no-code building blocks for your workflows. Quick Automate gives additional advanced options for power users such as control flow logic, case management, error handling, custom code, and more.

Getting started

- To get started with Quick Flows, see [the section called “Using Amazon Quick Flows”](#).
- To get started with Quick Automate, see [the section called “Using Amazon Quick Automate”](#).

Using Amazon Quick Flows

Amazon Quick Flows is a capability within Amazon Quick that lets any user create, customize, and share workflows that automate routine tasks. You can generate flows from conversations with chat agents, describe what you need in natural language, or build them manually using the visual editor — no technical skills required. Flows can also be published to an admin-managed library and shared with other Amazon Quick users in your organization.

Each flow is a sequence of steps that can gather user input, generate AI responses from your data or the web, take actions in connected applications, and apply logic to control how steps run.

What are Flows?

A flow is a sequence of *steps* that you define once and run whenever you need it. You build flows in the editor, then run them in either a guided step-by-step mode or a conversational chat mode where you can refine outputs and ask follow-up questions.

Each flow can include the following step types:

AI responses

- Chat agent step that gets a response from a custom agent and can take actions.
- Research step that conducts research on a topic.
- Web search step that gets a response using web search.
- General knowledge step that gets a response directly from base AI models.
- UI agent step to perform tasks on public websites.
- Create image step that generates an image from inputs.

Flow logic

- Reasoning group step that adds run instructions to one or more steps.

Data insights

- Quick data step that gets insights from spaces or knowledge bases.
- Dashboards and topics step that gets responses from dashboards and topics.

Actions

- Application actions step that reads or writes to connected apps.

User input

- Text step to get and use text from users.
- Files step to get and use files from users.

For detailed explanations of each component, see [the section called “Components and features”](#).

Why Flows?

Organizations face business processes that require both human judgment and system interactions. Amazon Quick Flows bridges this gap by combining AI reasoning with direct business actions, so you can automate tasks like generating responses for requests for proposals (RFPs), reviewing statements of work (SOWs), or compiling industry trends into a sales pitch.

Amazon Quick Flows can help across departments:

- *Sales, marketing, and operations*: Qualifying leads, generating personalized proposals, creating marketing content, updating customer relationship management (CRM) records, and supporting processes like RFP responses
- *Human resources (HR)*: Processing employee requests, answering policy questions, and automating onboarding steps
- *Finance*: Analyzing expense reports, flagging anomalies, and processing routine approvals
- *Information technology (IT)*: Automating troubleshooting, system monitoring, and access management

Getting started with Quick Flows

This section covers user personas, prerequisites, and the Amazon Quick Flows landing page interface.

User personas

Amazon Quick Flows serves three types of users:

Creator

Creators are business users who build flows. They understand the repetitive tasks and processes in their domain and translate them into reusable workflows.

End user

End users run flows that creators have built. They interact with flows to complete their daily tasks and provide feedback on how well the automation works.

Admin

Admins configure and manage governance controls for Amazon Quick Flows within Amazon Quick. They can:

- Enable or disable Quick Flows for the account
- Restrict access to Quick Flows for specific users
- Enable or disable Bedrock models for output refinement
- Require approval before users can share flows
- Unpublish shared flows and transfer ownership between users

Prerequisites for Quick Flows

Before you can create and use Amazon Quick Flows, you need to ensure that your Amazon Quick administrator has completed the following prerequisites.

Administrator setup requirements

Your Quick administrator must complete the following tasks before you can create and use Quick Flows:

- Set up and configure Quick for your organization. For more information, see [Setting up and signing in](#).
- (Optional) Restrict access to flows for specific users using custom permissions.

For browser and region requirements, see [the section called "Supported browsers"](#).

Required permissions

Permissions to create, run, share, and govern flows are determined by user subscriptions and any configured custom permissions. For more information, see [the section called “Managing subscriptions”](#) and [the section called “Managing user access”](#).

Amazon Bedrock model access

Quick Flows uses Amazon Bedrock models for AI reasoning in the General knowledge step. Your administrator must enable access in custom permissions for output refinement in flows using Bedrock models. For more information, see [the section called “General knowledge”](#).

Next steps

After ensuring that all prerequisites are met, you can:

- Learn about key concepts in Quick Flows. See [the section called “Key concepts”](#).
- Create your first flow. See [the section called “Creating flows”](#).

Amazon Quick Flows landing page

When you access Amazon Quick Flows, you are directed to the landing page. The landing page has three navigation tabs:

- **Flows library** — Browse, search, and manage your flows and flows that others have shared with you.
- **My schedules** — Monitor and manage scheduled flow runs.
- **Pending approval** — Review and act on flows that are awaiting approval.

Flows library

The Flows library organizes your flows into three views:

- **Recently used** — Flows you have recently opened or modified, sorted by most recent.
- **My Flows** — Flows you own or that have been shared with you directly.
- **All Flows** — Every flow available in your organization.

Use the search bar to find flows by name or description. For each flow, the context menu lets you open, duplicate, or delete it. Choose **Create flow** to start building a new flow. For more information, see [the section called "Creating flows"](#).

Using pre-built flows

Amazon Quick Flows provides pre-built flows and samples to help you get started.

Pre-built flows

Pre-built flows are ready-to-use templates available in **My Flows** on the Flows landing page. You can run them as-is or duplicate and customize them to fit your needs.

- **Blog post generator** — Creates blog content based on a topic, target audience, and key points you provide. The flow produces structured posts with headings, body content, and suggestions for tone and style.
- **Email composer** — Drafts professional emails for common business scenarios such as announcements, requests, and follow-ups. You specify the recipient context and purpose, and the flow generates appropriately toned content with a subject line.
- **Data summarizer** — Analyzes structured data such as CSV or tabular files and extracts key insights, trends, and anomalies. The flow presents findings in a summary format you can share or use in other steps.
- **Chart creator** — Generates visualizations based on your data and analysis goals. The flow recommends chart types based on your data and lets you configure the output for clarity.
- **Meeting summarizer** — Processes meeting notes or transcripts and extracts discussion topics, decisions, and action items. The flow organizes the output into a structured summary you can distribute to attendees.
- **Project status reporter** — Creates project status updates from key metrics and milestones you provide. The flow highlights risks, resource allocation, and next steps in a formatted report.

Flow samples

Flow samples are available on the **Create Flow** page. These are example prompts designed to help you construct your own flows. Each sample provides a detailed natural language description that you can use as a starting point.

For example, the "Find current industry trends and news" sample generates a prompt like:

"Help me prepare for my customer call by fetching latest news on the company. Accept customer name, industry, company size, meeting objective, and previous interaction notes as inputs. For the customer name and previous interaction notes, search from web to obtain key talking points and relevant industry trends. Finally, from the web analysis and previous interaction notes, generate a meeting brief that I can use as talking points along with relevant product recommendations from indexed data."

Customizing a pre-built flow

To customize a pre-built flow, duplicate it and modify the copy:

1. Select the pre-built flow you want to customize.
2. Choose **Duplicate** to create your own copy.
3. Open the duplicated flow in the Flow editor.
4. Modify the steps, prompts, and configuration as needed.
5. Save your customized flow.

For more information about editing flows, see [the section called "Editing flows"](#).

Flow components and features

These topics provide detailed information about individual step types and their capabilities. For configuration instructions, see [the section called "Editing flows"](#).

Terminology and key concepts

Understanding the core terminology and concepts helps you effectively create, run, share, and maintain flows within your organization.

Steps and @ references

A flow is made up of steps that each perform a specific function, such as calling an action, querying your data, or searching the web. Most steps are controlled through a natural language prompt.

Steps pass data to each other using @ references. When you write a prompt inside a step, type @ to see a menu of previous steps. Select one to include that step's output as context in your prompt. For example, if Step 1 collects a customer's issue and Step 3 needs to classify it, Step 3's prompt might say: "Classify the issue in @Customer Issue by severity (low, medium, high)."

The available step types are organized into the following groups:

AI responses

- **Chat agent** — Gets a response from a custom agent and can take actions in connected applications.
- **Research** — Invokes Amazon Quick Research to generate research reports as part of your workflow.
- **Web Search** — Generates responses using internet search results.
- **General knowledge** — Gets a response directly from Amazon Bedrock models, with configurable response preferences for speed or versatility.
- **UI Agent** — Navigates public websites that do not require a login and performs tasks like scrolling to find information or filling forms.
- **Create Image** — Generates AI images from text inputs.

Flow logic

- **Reasoning Group** — Groups related steps together with natural language instructions that define conditions, loops, validation, and execution order.

Data insights

- **Quick data** — Retrieves responses from spaces and knowledge bases.
- **Dashboards and topics** — Gets insights from Amazon Quick Sight dashboards and topics.

Actions

- **Application actions** — Performs read or write operations in connected third-party applications through pre-built connectors.

User input

- **Text** — Collects free-form text input from users.
- **Files** — Accepts file uploads from users for document processing.

Editor and Run mode

Editor mode is where you build your flow. You see all the steps laid out and can select each one to change its configuration. Run mode is where you test and execute your flow, with a chat panel where you can ask follow-up questions or refine the output.

AI response steps

AI response steps generate content using AI models. Amazon Quick Flows provides the following AI response step types.

Chat agent

Amazon Quick Flows allows you to use your chat agents to generate outputs from configured spaces or take action with configured action integrations, all within a workflow step.

Chat agents contain domain-specific knowledge, custom instructions, and connected tools. When you integrate a chat agent into a flow, you can automatically apply this specialized knowledge across multiple workflows without recreating it. For example, if you built a sales assistant chat agent that understands product details and follows brand guidelines, you can embed it in your outreach flow to ensure consistent communication at scale.

For configuration instructions, see [the section called “Editing flows”](#).

Note

The chat agent step is a single-turn interaction. The agent responds to the task you instruct it to do, but does not support a back-and-forth conversation within the same step.

Research

The research step invokes Amazon Quick Research to generate research reports within your flow. This lets you embed research directly into multi-step workflows — for example, creating account plans, conducting policy reviews, researching patent prior art, or generating industry reports.

For full details about Quick Research capabilities and limitations, see [Using Amazon Quick Research](#). For configuration instructions, see [the section called “Editing flows”](#).

You can reference the research output in later steps — for example, to send a summary over email to your team.

Web search

The web search step lets your flows retrieve current information from the internet. This is useful when you need to access real-time data, verify facts, or gather information from public sources beyond your organization's internal knowledge base.

Write a prompt describing what to search for. The search results can be referenced by later steps in your flow using @ references.

For configuration instructions, see [the section called “Editing flows”](#).

Note

Search results may vary over time as internet content changes. Some content may not be accessible through web search.

General knowledge

The General knowledge step generates text responses using Amazon Bedrock models. Instead of selecting a specific model, you choose a response preference, and Amazon Quick Flows automatically selects the most appropriate model based on your preference and the requirements of your flow.

Choose from:

- **Fast responses** — Optimized for speed across image, video, and text inputs.
- **Versatility and performance** — Balanced capabilities for diverse tasks.

Optionally adjust the creativity slider to control the randomness of the response.

If you do not see response preferences, verify that your administrator has enabled "Enable bedrock model usage in General knowledge step for output refinement" in the Custom Permissions page.

For configuration instructions, see [the section called "Editing flows"](#).

UI agent

The UI agent step (Preview) lets your flows interact with public websites that do not require a login. The agent can autonomously navigate websites, click, type, read data, and produce structured outputs — all described in natural language.

Writing effective instructions

- Be clear and specific about the task you want performed.
- Use single, complete URLs (for example, "Go to <https://example.com/reports>").
- Add constraints to narrow the scope (for example, "only look at the pricing section").
- Specify when the agent should stop (for example, "stop after finding the first matching result").
- Define the output format if needed (for example, "return the data as a bulleted list").

For configuration instructions, see [the section called “Editing flows”](#).

 Note

UI agent is currently in Preview. Some websites implement anti-automation measures such as CAPTCHA challenges that may limit UI agent capabilities. Websites that require login are not currently supported.

Create Image

The Create Image step generates AI images from text prompts. You can configure creativity level, exclude terms, and image seed in the advanced settings.

For configuration instructions, see [the section called “Editing flows”](#).

Flow logic steps

Flow logic steps control how your flow runs.

Reasoning Group

Reasoning groups give you control over how parts of your flow run using natural language instructions. A reasoning group contains its own set of steps — like an isolated workflow within your larger workflow — that runs based on conditions you define. You can add most step types to a reasoning group, except reasoning groups and research steps. Templates are available to help you get started.

Loops

You can repeat the steps in a group for each value in a list from a previous step's output. Reference the previous step in your instructions, and the Flows runtime handles the iteration for you. For example, if a previous step returns a list of customer emails, a reasoning group can process each email in turn.

Conditions

You can run the steps in a group based on natural language conditions that evaluate a previous step's output. For example, "Run if @Customer Priority is HIGH PRIORITY" routes only urgent items through the group's steps.

Validation

You can check inputs or outputs before proceeding. For example, a reasoning group can verify that a required field is present before passing data to an action step.

For configuration instructions, see [the section called “Editing flows”](#). For reasoning group limits, see [the section called “Quick Flows limits”](#).

Data insight steps

Data insight steps retrieve information from your Quick data sources.

Quick data

The Quick data step retrieves responses from your spaces and knowledge bases. Write a prompt describing what content to retrieve, and optionally link specific resources. By default, responses are generated from all knowledge sources the user has access to.

The system searches across indexed documents in your spaces, including knowledge bases from connected sources. For more information about setting up spaces and knowledge bases, see [Working with integrations](#).

For configuration instructions, see [the section called “Editing flows”](#).

Dashboards and topics

The Dashboards and topics step generates insights from your existing Amazon Quick Sight dashboards and topics. Responses can include charts, graphs, tables, and other visualizations.

Select a Quick Sight source (Dashboard or Topic) and write a prompt describing the insights you want. You can specify filters, date ranges, and other criteria in natural language.

For configuration instructions, see [the section called “Editing flows”](#).

Action steps

Action steps let your flows perform read or write operations in connected applications. The available operations depend on the action integrations that you have configured or that have been shared with you. For prerequisites, authentication methods, and available integrations, see [Working with integrations](#).

Action parameters

Some actions require or accept parameters such as filters or specific input values. To determine which parameters an action accepts, you can ask Quick — for example, "What parameters does list emails accept when filtering emails?"

Parameters can be:

- Supplied directly in the prompt (for example, "List emails with subject containing 'quarterly report'")
- Referenced from a previous step using @ references (for example, a user input step that collects a search term)
- Acquired dynamically during flow execution based on the flow context

Working with list results

When you use an action that returns a list of items (for example, listing emails or tickets), the results may not include every item or every detail. This is because many applications return results in batches rather than all at once, and Quick retrieves only the first batch to keep your flow fast and responsive.

If you need the full details for each item in a list, you can use a reasoning group to go through the results one at a time and retrieve the complete information for each. For example, you might list your open support tickets in one step, then use a reasoning group to get the full details of each ticket.

When doing this, be mindful of how many items you are processing. A large number of items means more steps to run, which increases the time your flow takes to complete and the amount of data in your results. Where possible, use filters in your initial list action to narrow down the results before processing them.

For configuration instructions, see [the section called “Editing flows”](#).

User input steps

User input steps collect information from users when they run a flow.

Text

The text input step collects free-form text from users. You can set a placeholder, a default value, and allow users to override the default at runtime.

Use placeholder text to guide users on what to enter. For example, you can present a set of options like "Enter 1 for Sales, 2 for Marketing, 3 for Support" to help users provide structured input that your flow can act on.

For configuration instructions, see [the section called "Editing flows"](#).

Files

The file upload step accepts a document, image, or video from users. You can upload a default file and allow users to override it at runtime. You can upload one file per step.

File uploads are subject to the same size and format restrictions as uploading files in chat. If your content exceeds these limits, consider using a space or knowledge base to process the request instead.

For configuration instructions, see [the section called "Editing flows"](#).

Creating and managing flows

This section covers how to create, edit, version, and share flows within your organization.

Creating flows

You can create flows using natural language, the visual editor, by duplicating an existing flow, or from a chat agent conversation.

Prerequisites for creating flows

Before creating a flow, ensure that you meet the requirements described in [the section called "Prerequisites"](#).

Creating a flow using natural language prompt (NLP)

The natural language prompt method lets you describe your desired flow in plain language, and Amazon Quick Flows generates a flow based on your description.

1. Sign in to the Quick console.
2. In the navigation pane, choose **Flows**.
3. Choose **Create Flow**. You see a page where you can enter a natural language prompt or use a sample prompt.
4. In the prompt field, describe the flow you want to create. Be as specific as possible about:
 - The purpose of the flow

- The inputs it should accept
- The processing it should perform
- The outputs it should produce
- The actions it should take

5. Choose **Generate Flow**.

6. Review the generated flow and make any necessary adjustments.

7. Switch to **Run mode** to test your flow. If the flow is ready for use, you can share or publish it.

 Note

The quality of the generated flow depends on the clarity and specificity of your prompt. You may need to refine your prompt or make manual adjustments to achieve your desired outcome.

Creating a flow from scratch (blank flow)

Creating a flow from scratch gives you complete control over the design and functionality of your flow.

1. Sign in to the Quick web experience.
2. In the navigation pane, choose **Flows**.
3. Choose **Create Flow**.
4. Create a blank flow.
5. Enter a name and optional description for your flow.
6. In the Flow builder, add and configure the steps you need:
 - Add user input steps to collect text or files from users.
 - Add AI response steps to generate content from chat agents, web search, general knowledge, or image generation.
 - Add data insight steps to get responses from Quick data or dashboards and topics.
 - Add reasoning groups to control how steps run with conditions, loops, or validation.
 - Add action steps to read or write to connected applications.
7. Choose **Save** to save your flow.

For more information about the different components you can use in a flow, see [the section called "Key concepts"](#).

Creating a flow by duplicating an existing flow

Duplicating a flow creates a copy that you can modify for your own purposes.

1. Sign in to the Quick console.
2. In the navigation pane, choose **Flows**.
3. Find the flow you want to duplicate.
4. Choose the ellipsis (:) next to the flow name, then choose **Duplicate**.
5. Modify the name and optionally the description for your purposes.
6. Make any necessary modifications to the flow.
7. Choose **Save** to save your flow.

Note

When you duplicate a flow, all components, connections, and configurations are copied. However, you may need to reconfigure certain settings, such as authentication for connectors.

Creating a flow from an agent conversation

You can create a flow from a conversation with a chat agent. For example, if you have been working through a task like reading an email, augmenting it with internal or web knowledge, and drafting a response, you can convert that conversation into a reusable flow.

1. Sign in to the Quick console.
2. Start a conversation with a chat agent.
3. Work through the task you want to automate.
4. Choose the Quick Flows icon in the chat bar.
5. Choose **Create Flow from conversation**. Amazon Quick Flows interprets the interactions in your conversation and brings you to the Create flow screen with a generated natural language prompt, name, and description.
6. Review and edit the prompt, name, and description as needed.

7. Choose **Generate Flow**.
8. Review the generated flow and make any necessary adjustments.
9. Choose **Save** to save your flow.

Flow creation requirements

When creating a flow, keep the following requirements in mind:

Flow name

Flow names must be unique within your Quick account and can contain up to 128 characters.

Flow components

A valid flow must have at least one output step or action step.

Reasoning groups

Each reasoning group must have at least one instruction that defines what the AI model should do with the inputs.

Action steps

Action steps require proper authentication and configuration before they can be used in a flow.

Next steps

After creating a flow, you can:

- Test your flow to ensure it works as expected.
- Publish your flow to make it available to users. See [the section called "Publishing changes"](#).
- Share your flow with specific users or groups. See [the section called "Sharing flows"](#).
- Monitor your flow's usage and performance. See [the section called "Progress tracker"](#).

Editing flows

After creating a flow, you can edit and configure it in the Flow editor.

Accessing the Flow editor

1. Sign in to the Amazon Quick console.

2. In the navigation pane, choose **Flows**.
3. Find the flow you want to edit.
4. Choose the flow tile, or choose the ellipsis (:) and select **Open**.

Configuring step types

Most steps with a prompt field support @ references to include data from previous steps. See each step's configuration for specific limitations.

Configuring text input steps

In Editor mode, add or select a text input step. In the configuration panel, set the following:

- **Title:** A name for the step.
- **Placeholder:** Optional text that appears inside the input field when it is empty. This text is not used for the flow run.
- **Default Value:** Optional. The value provided if the user doesn't enter an input.
- **Allow override of default value:** Toggle to let users replace the default value at runtime.

Configuring file upload steps

In Editor mode, add or select a file upload step. In the configuration panel, set the following:

- **Title:** A name for the step.
- **Default file:** Optional. Upload a document, image, or video to use if the user doesn't provide one. You can upload one file per step.
- **Allow override of default value:** Toggle to let users replace the default file at runtime.

Configuring general knowledge steps

In Editor mode, add or select a general knowledge step. In the configuration panel, set the following:

- **Output preference:** Choose one:
 - **Fast responses** — Across image, video, and text inputs.
 - **Versatility and performance** — Balanced capabilities for diverse tasks.

- **Prompt:** Write the prompt that instructs the model what to generate.
- **Creativity Level:** Optional. Adjust the slider from low to high to control the randomness of the response.

For more information, see [the section called “General knowledge”](#).

Configuring Quick data steps

In Editor mode, add or select a Quick data step. In the configuration panel, set the following:

- **Prompt:** Write the prompt that instructs what content to retrieve.
- **Link specific resources:** Select spaces and knowledge bases to get insights from. By default, responses are generated from all knowledge sources the user has access to.

Configuring web search steps

In Editor mode, add or select a web search step. In the configuration panel, set the following:

- **Prompt:** Write the prompt that instructs what content to search for on the web.

Configuring research steps

In Editor mode, add or select a research step. In the configuration panel, set the following:

- **Title:** A name for the step.
- **Research objective:** Describe what you want to research.
- **File uploads:** Optional. Upload default files to help guide your research.
- **Research materials:**
 - **Preferred websites:** Optional. Specify websites or types of websites the agent should prioritize (for example, government websites, academic journals).
 - **Websites to avoid:** Optional. Specify websites or types of websites to exclude (for example, social media, blogs).
- **Data and apps:** Select all data and apps, or choose specific ones.

Configuring chat agent steps

In Editor mode, add or select a chat agent step. In the configuration panel, set the following:

- **Title:** A name for the step.
- **Chat agent:** Select the agent to use.
- **Prompt instruction:** Write the prompt that instructs the agent.
- **Data and apps:** Optionally narrow down the selected data and apps to refine your use case.
- **Web search:** Toggle to enable or disable web search for the agent.

Configuring UI Agent steps

In Editor mode, add or select a UI Agent step. In the configuration panel, set the following:

- **Title:** A name for the step.
- **UI Agent Instructions:** Write the instructions for the UI agent. Use single, complete URLs for faster, more accurate results.

Configuring Create Image steps

In Editor mode, add or select a Create Image step. In the configuration panel, set the following:

- **Prompt:** Describe the image to generate.
- **Advanced settings:** Configure creativity level, exclude terms, and image seed.

Configuring Dashboards and topics steps

In Editor mode, add or select a Dashboards and topics step. In the configuration panel, set the following:

- **Quick Sight source:** Choose from Dashboard or Topic.
- **Prompt:** Describe the insights you want from your dashboard or topic.

For more information, see [the section called “Dashboards and topics”](#).

Configuring action steps

In Editor mode, add or select an action step. In the configuration panel, set the following:

- **Connector:** Select the connector to use (for example, Salesforce, Jira, Slack).
- **Action:** Select the specific action to perform.

- **Prompt:** Write prompt instructions to execute your actions.

For more information, see [the section called "Actions"](#).

Configuring reasoning groups

In Editor mode, add or select a reasoning group. In the configuration panel, set the following:

- **Instructions:** Write instructions that tell the model what to do, such as conditions, loops, or validation logic. You can reference the output of a previous step as the input to a loop or conditional statement.

Prompt writing for steps

When writing prompts for steps, consider the following best practices:

- Be clear and specific about what you want the model to generate.
- Provide context to help the model understand the task.
- Specify the desired format, tone, and style of the output.
- Use specific language to control the output — for example, "respond in bullet points", "limit the response to three sentences", or "use a formal tone".
- Use examples to illustrate the expected output when appropriate.

Adding and removing steps

Adding steps

1. In the Flow editor, choose **Add step** from the toolbar.
2. Select the type of step you want to add.
3. Drag the step to the desired position in your flow.
4. Configure the step as needed.

Removing steps

1. Select the step you want to remove.
2. Choose **Delete** from the context menu.
3. Confirm the deletion when prompted.

 Note

When you remove a step, any @ references to that step in other steps are also removed. You may need to update other steps to maintain your flow.

Publishing changes

After making changes to your flow, publish them to make them available to users.

1. In the Flow editor, choose **Save** to save your changes.
2. Choose **Publish**.
3. Choose **Publish** to confirm.

When you publish, the changes become immediately available to all users who have access to the flow.

Updating Flow details

Updating title and description

1. In the Flow editor, directly edit the title and description in-line.
2. Choose **Save**.
3. Choose **Publish**.

Viewing the original prompt

If your flow was created using a natural language prompt, you can view the original prompt:

1. In the Flow editor, select the kebab menu in the header.
2. Choose **View prompt**.

 Note

The original prompt is read-only. To create a new flow based on a modified prompt, use the natural language prompt method.

Best practices for editing flows

- Test your changes thoroughly before publishing them to ensure they work as expected.
- Use clear and descriptive names for steps to make your flow easier to understand and maintain.
- Write detailed prompts and instructions to get the best results from AI models.
- Use @ references to create dynamic flows that adapt to user inputs.
- Consider the user experience when designing the flow of steps.
- Document your changes in the publication description to maintain a clear history of updates.

Step configurations and available features may change over time. For the latest information, see the [the section called “Key concepts”](#) and the [Amazon Quick Flows product page](#).

Versioning

Amazon Quick Flows uses versioning to let you edit a flow without affecting the version that others are using. You can have up to three versions of a flow at any time.

How versioning works

Your organization's setup determines which version types you see.

If your organization uses approval review

You work with three version types:

- **Draft** — Your working copy. You can edit and test your flow without affecting the published version. As you edit, your flow is saved automatically. You can have one draft at a time.
- **Pending Approval** — A version submitted for review. You can view it and check the approval status, but you cannot edit it. If changes are requested, update your draft and resubmit.
- **Published** — The approved version that other users can access and run.

For more information about approval reviews, see [the section called “Approval review”](#).

If your organization doesn't use approval review

You work with two version types:

- **Draft** — Your working copy. You can edit and test your flow without affecting the published version. As you edit, your flow is saved automatically. You can have one draft at a time.

- **Published** — The version that other users can access and run. You can publish directly when your flow is ready.

Version Type	Can Edit?	Visible to Users?	Next Step
Draft	Yes	No	Publish (or submit for approval)
Pending Approval*	No	No	Wait for approval
Published	No	Yes	Create new draft for changes

*Pending Approval only applies if your organization uses approval review.

Publishing changes

Any changes to a flow's title, description, steps, logic, or sharing permissions require publishing. When you publish, the new version replaces the current published version for all users. End users always see only the latest published version.

If your organization uses approval review, publishing submits your flow for review. Once approved, it becomes the new published version. Without approval review, publishing makes changes available immediately.

Unpublishing flows

You can unpublish a shared flow to remove it from general availability while keeping it accessible to co-owners. Viewers lose access, and the flow returns to draft state. Unpublishing does not require approval review.

Sharing flows

When you share a flow, you give others permission to view and run it, or to collaborate on it as co-owners.

Sharing options

You can share a flow in three ways:

- **With individuals** — Share with specific people using their email addresses.
- **With groups** — Share with pre-defined groups in your organization (for example, departments or project teams).
- **With everyone** — Make the flow available to all users in your Amazon Quick Flows instance.

Depending on your organization's settings, sharing may require approval review. For more information, see [the section called "Approval review"](#).

How to share a flow

1. Open your flow.
2. Choose **Share**.
3. Select how you want to share: with individuals, groups, or everyone.
4. For individuals, enter email addresses. For groups, search for and select the group.
5. Choose the permission level: **Viewer** or **Co-owner**.
6. Choose **Share**.

Permission levels

Viewer

Viewers can run the flow and view results, but cannot edit, share, or delete it. Viewers only see the published version.

Co-owner

Co-owners can edit the flow, share it with others, manage permissions, publish changes, and access draft versions. Co-owners cannot remove the original owner. Only one person can edit at a time.

Permission reference

Permission	Can run?	Can edit?	Can share?	Can delete?
Viewer	Yes	No	No	No
Co-owner	Yes	Yes	Yes	Yes

Approval review

When approval review is enabled, users must submit their flows for review before they can be shared. This gives administrators oversight of what flows are available in the organization.

How approvals work

When approval review is enabled, sharing a flow follows a three-stage cycle:

- **Submitted** — The flow is waiting for review. You can view it and check the status, but cannot edit it. You can withdraw the request if needed.
- **Rejected** — The flow was not approved. You can view the reviewer's feedback, update your draft, and resubmit.
- **Approved** — The flow is shared with your intended audience. Any future changes require a new approval cycle.

Who can review flows

Users with Amazon Quick Enterprise subscriptions (Author Pro or Admin Pro roles) can review and approve flows. Users with Amazon Quick Professional subscriptions cannot review flows.

Reviewers can:

- View all submitted flows in the **Pending approval** tab
- Approve or reject flows individually or in bulk
- Provide feedback when rejecting a flow

Submitting a flow for approval

1. Complete and test your flow.
2. Choose **Share** and specify who should have access.
3. Choose **Submit for review**.
4. Monitor the approval status in the **Pending approval** tab.

If your flow is rejected, review the feedback, update your draft, and resubmit.

Enabling and disabling approval review

Administrators control approval review through the Custom Permissions page.

When enabled:

- All future flow sharing requires approval.
- Existing shared flows remain accessible.
- Users can still create and edit flows in draft mode.

When disabled:

- All pending approval requests are automatically rejected.
- Users can share flows immediately without approval.
- Previously approved flows remain shared.

Approval states reference

State	Can edit?	Visible to users?	Available actions
Draft	Yes	No	Submit for approval
Submitted	No	No	View status, withdraw request
Rejected	Yes	No	Update and resubmit
Approved	No	Yes	Create new draft for changes

Running flows

These topics cover how to run flows in different runtime modes, download outputs, and schedule recurring runs.

Interacting with flows in runtime mode

When you open a flow, you select **Run mode** to execute it. There are three ways to start a flow.

Starting a flow with structured input

If the flow has input steps, enter the required information and choose **Start**. The flow runs through each step in sequence and displays the results.

Starting a flow with conversational runtime

You can supply the flow input (optional) and ask the agent to start the flow through the chat panel. The agent collects any additional inputs, runs the steps, and presents results within the conversation. You can ask follow-up questions or refine outputs as the flow runs.

Starting a flow from a chat agent

You can run a flow directly from a chat agent conversation:

1. Choose the Flows menu in the chat footer.
2. Select the flow you want to run.
3. Follow the prompts to execute it.
4. When finished, choose **End** to return to your chat agent conversation.

Details from your conversation can automatically populate the text input to your flow. When the flow completes, the output is shared with the calling agent.

Progress tracker

The progress tracker shows the status of each step as the flow runs.

Resuming runs from history

You can view and resume previous flow runs from the flow's history. Open the flow and choose the history icon to view previous runs.

Downloading output in Amazon Quick Flows

After a flow run completes, you can download one or more step outputs as a Word (.docx) or PDF document. You can save files locally or to a shared Quick space for team collaboration. Downloaded files maintain the original formatting, including markdown structure, visual elements, and data tables.

Download flow outputs

1. In Run mode, choose the download icon (downward arrow) in the top right.

2. Choose **Download as**.
3. Select which outputs to include:
 - Choose **Select All** to include all steps in chronological order.
 - Use checkboxes to select specific steps.

 Note

User input and file upload steps cannot be downloaded.

4. Select your preferred file format:
 - **Word (docx)** — Editable document format
 - **PDF** — Fixed-layout document format
5. Choose your download destination:
 - **Download** to save the file locally.
 - **Save to a space** to save to a shared space (requires edit access).

Scheduling your Amazon Quick Flows

Schedules let you automate recurring flow runs without manual intervention. You can schedule flows you have created or flows that have been shared with you. All schedules are private to you and cannot be shared with other users.

Common uses include generating recurring reports, summarizing open items from external services, or preparing daily meeting briefings.

Creating a schedule

1. Open your flow in Run mode.
2. Choose the scheduling icon.
3. Choose **Create schedule**.
4. Configure your schedule:
 - **Schedule name** — A unique name for the schedule.
 - **Description** — Optional. Details about the schedule's purpose.
 - **Repeat configuration** — Choose from suggestions or configure custom recurrence (daily, weekly, or monthly) with start date, optional end date, and time zone.

5. Choose **Next**.
6. Provide default inputs for each scheduled run.
7. Choose **Next**.
8. Configure action permissions:
 - Turn on **Run with no confirmation** to automatically submit action forms during scheduled runs.
 - Turn off to review and confirm each action manually.
9. Choose **Save**.

Note

Action permissions only appear if the flow contains action steps. Automatic form submission is enabled by default for write actions. User confirmation is recommended to help avoid AI prediction errors.

Managing schedules

You can manage schedules from two places:

- **From the flow** — Open the flow, choose the scheduling icon, and edit, pause, duplicate, view runs, or delete a schedule.
- **From the Flows library** — Choose the **My schedules** tab to see all your schedules across all flows with their status, frequency, last run, and associated flow name.

Schedule history

All manual and scheduled runs appear in the flow run history. You can filter by all runs, scheduled runs, or failed runs. History is retained for 30 days.

Visual indicators help distinguish run types:

- Schedule run tag for scheduled runs
- Orange dot for un-viewed completed runs
- Running tag for in-progress runs
- Failed tag for failed runs

Note

You cannot view an ongoing scheduled run until it completes.

Notifications

You receive email notifications when:

- A scheduled run completes successfully (with a link to the results)
- A run requires your input to continue
- A run requires authentication for a third-party application
- A flow with your schedules is unshared from you (schedules are archived)
- A flow with your schedules is deleted
- A flow with your schedules is updated with a new version (schedules continue with the updated version)

Authentication for connectors

Ensure your authentication for connectors is current before scheduling. If authentication expires before a scheduled run, you receive an email notification. To verify, run the flow manually or visit the connector page in Quick to sign in.

Managing Quick Flows

As an administrator, you can manage how Amazon Quick Flows operates within your organization, including governance policies, user permissions, and feature availability.

Administrative access requirements

To manage Quick Flows settings, you must have administrator privileges in your Amazon Quick instance. Administrative controls are available through the Amazon Quick console's management interface, where you can configure permissions and manage assets.

Governance capabilities overview

Quick Flows administrative controls provide multiple layers of governance:

- **Feature enablement** - Control whether flows are available to users in your account

- **Approval workflows** - Require administrative approval before flows can be shared with users
- **Asset management** - Oversee Flow sharing, ownership, and visibility across your organization

Managing Quick Flows permissions and assets

Administrators manage Quick Flows permissions and assets through two interfaces in the Quick console: the Custom Permissions page and the Asset Management page.

Custom Permissions page

The Custom Permissions page controls which Quick Flows capabilities and features are available to users.

Restrict capabilities

- **Enable or disable flows** — Control whether Quick Flows is available to users in your account. When disabled, all Quick Flows functionality is unavailable, but existing flows are preserved. Quick Flows is enabled by default for new Quick instances.

Restrict features

- **Allow creators to share without approval** — When enabled, creators can share flows without review. When disabled, all eligible users can review and approve flow sharing requests.
- **Enable bedrock model usage in General knowledge step for output refinement** — Controls whether Bedrock models are available for General knowledge steps. Enabled by default.
- **Enable UI agent to perform browser tasks (Preview)** — Controls whether UI agent steps can be used in flows.

Asset Management page

The Asset Management page provides oversight of all flows in your organization.

- **Share flows** — Share flows on behalf of creators, redistribute flows to different groups, or manage sharing permissions across the organization.
- **Transfer ownership** — Transfer flow ownership between users, such as when a creator leaves the organization.
- **Unlist flows** — Remove flows from the shared library, making them unavailable to viewers while preserving them for the creator to edit. Unlisted flows can be re-shared once issues are resolved.

Quick Flows limits

Flow limits

The following limits apply to flow creation and management in your Quick instance:

- **Maximum flows per instance:** 10,000 flows
- **Maximum flows per user:** 100 flows
- **Maximum steps per flow:** 35 steps

Input limits

Input limits vary depending on whether you are using General knowledge or Quick data as your data source.

Text input limits

- **General knowledge:** 40,000 characters
- **Quick data:** 15,000 characters

File upload limits

File upload limits depend on your knowledge source and file type:

File size limits:

- **General knowledge:**
 - Document files: Up to 50 MB (supported file types depend on output preference)
 - Video files: Up to 1 GB
 - Image files: Up to 4.5 MB
- **Quick data:**
 - Document files: Up to 50 MB (.pdf, .txt, .rtf, .doc, .docx, .ppt, .pptx, .csv, .xls, .xlsx)
 - Image files: Up to 10 MB (.png, .jpg, .jpeg)

Maximum number of files:

- **General knowledge:** Up to 5 files

- **Quick data:** Up to 20 files

Character limits for file content:

- **General knowledge:**
 - Faster responses: Up to 1M characters total context limit
 - Versatility and Performance: Up to 1M characters total context limit

The total context limit includes characters from output prompts defined at build time and characters from input steps and file upload steps at runtime.

- **Quick data:** 665,000 characters

Output limits

Output limits control the size of prompts and generated responses:

- **Maximum input prompt size:**
 - General knowledge: 5,000 characters
 - Quick data: 5,000 characters
- **Maximum output characters rendered:**
 - General knowledge: Up to 40,000 characters
 - Quick data: 8,000 characters

Reasoning group limits

- **Maximum iterations per reasoning group:** 50 iterations

When referencing the output of a reasoning group, especially looped content, the output is bound by the input character limits of the receiving step. For larger outputs, consider staging the results and querying them in a later step.

Schedule limits

- **Region availability:** Schedules in flows are currently supported in US East (N. Virginia), US West (Oregon), and Europe (Ireland)

- **Maximum schedules per user:** 20 schedules per user
- **Maximum schedules per instance:** 10,000 schedules

Regional availability

Certain features have regional limitations:

- **Image generation:** Currently supported in US East (N. Virginia), US West (Oregon), and Europe (Ireland)

Using Amazon Quick Automate

Overview

Amazon Quick Automate is an AI-powered application that creates sophisticated automations using natural language or documentation. Amazon Quick Automate revolutionizes enterprise workflow by transforming complex processes into intelligent, adaptive automations.

Note

Amazon Quick Automate is available only for Amazon Quick accounts provisioned through the AWS Management Console. It is not available for Free or Plus accounts created at aws.com/quick.

Enterprises struggle to automate end-to-end business processes due to the growing complexity of their application ecosystems and the limitations of traditional automation tools. Amazon Quick Automate provides the easiest way for businesses to build, deploy, and continuously maintain workflow automations with the set of capabilities needed to support complex use cases that span multiple applications, data sources, and departments.

Customers can simply provide their process requirements by uploading process documents or simply describing their business process in a natural language. From the documents or description, Amazon Quick Automate creates a plan with steps and actions for your automation. You can expand upon the plan by adding additional steps or actions or modify the existing steps or actions. You can add additional steps by dragging actions or steps from an extensive Actions library.

Amazon Quick Automate's intelligent agents can seamlessly navigate and interact with web applications, make contextual decisions based on real-time data, execute API calls, and even generate custom code to handle complex scenarios. When human oversight is required, the system can create human-in-the-loop tasks for manual review or approval.

Key concepts

Automation groups

Automation groups create logical groupings for automations, allowing organizations to manage permissions and resources efficiently across different departments or projects. These groups provide a way to organize your automations and setup permissions based on business units, functional areas, or project teams.

By organizing automations into groups, you can apply consistent access controls, resource restrictions, and governance policies. This organizational structure supports scalable automation management as your organization's automation needs grow and evolve.

Project

A *project* is the primary container that encapsulates all metadata, assets, and contextual information related to an automation initiative. It serves as the foundational unit for organizing, authoring, testing, deploying, and maintaining automations across their lifecycle.

Lifecycle of a project

- **Create a project:** Define the automation's scope including its name, purpose, business case, and attach process documentation.
- **Build automation:** Within the project, build and test and automation.
- **Version and deploy:** Publish automation versions for testing or production. Projects maintain version history and deployment status.
- **Monitor and iterate:** Review logs, metrics, for each run and iterate on the automation. Each iteration creates traceable project-level change history.

Automations

An *automation* is a visual artifact that represents the executable logic of a process. It is created using the *Visual Designer* within , where you define the sequence of actions, decisions, and agent interactions needed to complete a business process. Each automation serves as the operational blueprint for how work is performed, combining AI-powered agents, system actions, and human inputs in a unified flow.

Actions

Actions are the building blocks for your automations. Each action performs a specific task, from reading data to interacting with external systems.

Agents

An *agent* is an AI-powered action that can understand natural language input, reason through scenarios, and perform complex tasks on your behalf. Agents bring intelligence and adaptability to automations by interpreting intent, handling variability, and executing actions dynamically.

- **UI agent** is a native agent that understands natural language instructions to perform complex browser actions. It can autonomously navigate websites, click, type, read data, and produce structured outputs optimized for downstream automation steps.
- **Custom agent** is an AI-powered action that processes natural language inputs to automate complex steps using integrated tool-calling capabilities. It primarily uses integrations as its tool interface, while offering extensibility to use Code as tool, and other native actions like human-in-the-loop task. The agent delivers structured, predictable outputs optimized for seamless integration into downstream automation steps.

Getting started

Amazon Quick Automate is a browser-based solution that requires no installation, has minimal IT dependencies, and is always up to date. Amazon Quick Automate is designed for **technical teams and power users** who build, test, and maintain intelligent automations across enterprise systems. These users typically have a strong understanding of business processes, system integrations, but may not necessarily be software developers. Amazon Quick Automate equips advanced users and technical teams with the flexibility, control, and depth needed to design and scale enterprise-grade automations. To get started with Amazon Quick Automate, simply access it through your Amazon Quick and begin exploring the automation capabilities that best fit your needs.

Key capabilities

- Multi-agent architecture
 - Planning Agent designs automations from natural language inputs
 - UI Agent handles interface interactions
 - Specialized agents for process-specific tasks
 - Integration with

- Enterprise controls and governance
 - Granular access control management
 - Comprehensive activity logging
 - Built-in version control
 - Governance features for compliance
- Human-in-the-loop capabilities
- Case management
- Exception handling support
- Real-time monitoring
- Built in scalability

Setup tasks

Complete these setup tasks to prepare your Amazon Quick Automate environment for creating automations.

Sign up

Before you can use Amazon Quick Automate, you need to sign up for and create your account. The sign-up process establishes your initial workspace and provides access to Amazon Quick Automate.

Set up automation groups

Automation groups create logical groupings for automations, allowing organizations to manage permissions and resources efficiently across different departments or projects. These groups provide a way to organize your automations and setup permissions based on business units, functional areas, or project teams. By organizing automations into groups, you can apply consistent access controls, resource restrictions, and governance policies. This organizational structure supports scalable automation management as your organization's automation needs grow and evolve. After signing up, configure automation groups and permissions to organize your automations and control access across your organization. This step ensures that different departments and teams can work with automations appropriate to their roles and responsibilities.

Setting up proper automation groups and permissions from the beginning helps maintain security and compliance as your automation environment grows. You'll establish the organizational structure that supports efficient workflow management across multiple departments or projects.

Setup user roles

Amazon Quick Automate provides granular role based access: Viewer, Contributor, and Owner. Viewers can submit new project ideas, interact with human in the loop tasks, and access monitoring dashboards. Contributors can build, test, and deploy automations. Owners can control automation access to systems through integrations and secure credentials as well as user access.

Configure integrations

- After setting up AWS service integrations, you can begin using them in Quick Automate. See [the section called “Connectors”](#) for setup details.
- First, add the integration to an Automation Group. This allows you to control access for which automations can use this integration. To add an integration to your Automation Group:
 - Go to the **Projects** page
 - Click **Manage groups** and select the group you want to add an integration to
 - In the **Assets** sections, click **Add**, and then **Actions**
 - Select the Integration you created earlier and click **Add** to link it to your Automation Group
- Now when users create projects in that Automation Group, they will find the actions associated with this Integration and be able to select this Integration for authentication

Set up Virtual Private Cloud (VPC) connections

Amazon Quick Automate can access publicly hosted internet endpoints by default. For more complex setups, including accessing privately hosted websites, Amazon Quick Automate automation jobs can optionally be associated with your VPC. For more information, see [the section called “Setting up a VPC to use with Amazon Quick”](#).

With this configuration in place, all outbound network traffic from the managed browser is routed through an ENI in your configured VPC. From there, you can send that traffic to privately hosted resources or use it to obtain a persistent IP for IP allow listing.

The VPC-connected Amazon Quick Automate setup provides:

- **Direct access** to privately hosted resources in the connected VPC
- **Extended access** to other privately hosted resources in peered, PrivateLinked, etc. VPCs
- **Internet access** through your managed Internet Gateway

Setup instructions

Step 1: Choose your VPC

Select an existing VPC or create a new one that includes private and public subnets in one or more Availability Zones, NAT and Internet Gateways.

Note

If creating a new VPC, the default selections in AWS VPC creation wizards are sufficient.

Step 2: Setup connectivity to target resources

Bot traffic will egress through an ENI in your connected VPC. Choose the appropriate connectivity method:

For resources in the same VPC:

- No additional setup required

For resources in different VPCs, use one of these AWS constructs:

- AWS VPC Endpoint Services (PrivateLink) - Exposing individual service endpoints from one VPC to another
- AWS Transit Gateway - Centralized management of a cluster of peered networks
- AWS VPC Peering - Direct connection between a pair of VPCs

Step 3: Create Quick VPC Connection

Use Quick Admin tools to create a VPC Connection. This step creates two ENIs in the private subnet within your chosen VPC.

Step 4: Associate VPC Connection with Automation Group

Once the VPC Connection is created, associate it with your automation group from the VPC connection section in the automation group page. After this association is configured, all subsequent automations run within the specified group will use the VPC Connection for network access.

Create project

This section walks you through the process of creating a new automation project in Amazon Quick Automate.

To create a new project

1. Click the **Create Project** button to begin.
2. Fill in the project details in the Provide project details section:
 - **Name:** Enter a descriptive name for your automation project.
 - **Group:** Select the appropriate group for your automation.
 - **Description** (optional): Provide additional context about the project's purpose.
 - **Upload Documents** (optional): Attach existing documentation to jumpstart your automation.
3. Click **Next: Business case** to proceed to the business case section.

Defining the business case (Section is optional)

This section is optional but recommended for tracking ROI and prioritization. Complete the following fields:

- **Hours saved per case:** Estimate time savings per instance.
- **Cases per year:** Enter the expected annual volume.
- **Hours saved per year:** This may calculate automatically based on previous entries.
- **Project priority:** Set the relative importance of this automation.
- **Target launch date:** Define the target implementation date.

Click **Create** to finalize and generate your new automation project.

Note

Required fields are marked accordingly in the interface. Optional fields provide additional context but can be completed later if needed.

After creating your project, you can access a summary page with Overview, launch goals and document section will open for review. This comprehensive layout allows you to quickly grasp the essential information about your project or resource. You will have Summary (default selected), Versions and Deployments tabs.

In Summary, the Overview section provides a high-level description and launch Goals. The Documents section contains relevant document uploaded. On the right-hand side of the page, you'll notice a Status panel. This panel offers real-time updates on the current state of your project. It may include information such as deployment status, health checks, or any ongoing processes.

Building automations

Amazon Quick Automate provides multiple methods for creating automations to suit different needs and skill levels. This section describes the available creation methods and tools.

Prerequisites

Before creating an automation, you need to create a project within a specific automation group. The automation group will control what integrations and credentials are available to you when building that automation. Ensure the required integrations and credentials are set up for any external systems you plan to interact with in your automation.

Once you have created a project, click the **Start building** button to get started.

Note

You must have owner permissions to manage integrations and credentials for an automation group.

Creation methods

Choose from three approaches to build your automation. If you're new to Amazon Quick Automate, we recommend:

- Start by exploring sample automations
- Try creating a simple automation using natural language
- Experiment with manual editing once you're familiar with the basics

Using sample automations

Get started quickly by exploring pre-built samples that demonstrate common automation patterns. These samples provide hands-on examples of different automation capabilities and serve as learning resources to help you become familiar with the capabilities of Amazon Quick Automate. To use a sample automation:

- Search to explore different samples by industry.
- Click on the sample to view an overview of the automation.
 - If any integrations are required, they will be listed as prerequisites.
- Click **Start with sample** to open the automation in the canvas.
- You will now be able to view the automation, edit it, and test it.

Creating with natural language

The Automation Assistant provides an AI-powered approach to creating automations through:

- Chat interface - Describe your process conversationally in natural language
- Document upload - Import any documentation that describes your process steps and requirements

When you provide input through either method, the Automation Assistant will generate your automation through a two-step guided experience:

- **High-level plan**
 - First, the Assistant analyzes your requirements and generates a high-level automation plan. The plan consists of **process steps** that logically group the related actions as part of the automation. Each process step initially contains natural language instructions.
 - Review and edit the instructions for each process step and then click **Generate** to create the low-level actions for that step.
- **Low-level actions**
 - Next, the Assistant will build the low-level actions and process logic based on the step instructions.
 - Review and edit the generated actions. Click **Run** or **Debug** to begin testing.

Best practices for providing natural language inputs:

- Specify the step by step process as if you were training a new hire
- Include exact details needed to perform each step:
 - Email address if sending a message
 - File name and location if uploading/downloading a document
 - URL if navigating to a website

Creating from scratch (blank automation)

Build automations manually using the visual designer interface:

- Navigate to your project summary
- Click **Start building**
- Select **Skip** to access the canvas

The designer interface provides several key components:

- **Canvas** - Visualizes the automation. You can zoom in and out, expand and collapse process steps, and reorder actions on the canvas.
- **Actions panel** - The actions panel has a comprehensive list of all the automation actions made available within your automation group. You can search and filter to find actions organized by category (Agents, Process flow, Web browser, etc.). Drag-and-drop actions onto the canvas to build your automation. You can also click the plus sign found when hovering between actions to add a new action directly in place on the canvas.
- **Properties panel** - Once you add an action or click on an existing one in your process, the properties panel allows you to configure the input and output parameters that control how that action will behave.
- **Settings panel** - Allows you to create and edit runtime configurations and explore the credentials available to your automation.
- **Variables panel** - Displays all the variables used within your automation. Creating a variable is seamless by adding a new reference directly in your automation.

Runtime configurations and variables

Runtime configurations allow you to create easily editable settings that are used by your automation. Examples of runtime configurations include:

- Website URLs
- File names and locations
- Email configurations like subjects and sender lists

Variables are used to store and pass information between the actions in your automation when it runs. Add a new variable to your automation by simply providing a new reference name in an output property. Common variable types include:

- String - Plain text (e.g., name, description)
- Number - Numeric value (e.g., quantity, score)
- Boolean - True/False value
- Array - Collection or list of items
- Object - Key-value pairs
- File - Documents and media (e.g., PDFs, images)
- Data table - Spreadsheet-like data with rows and columns

Editing automations

Once you've created your automation, you can modify it using either the Automation Assistant or by directly editing on the canvas.

Using Automation Assistant

The Automation Assistant helps you make changes through natural language. You can:

- Ask for updates to the entire automation, specific steps or specific actions
 - Focus the chat on specific steps by selecting them first
- Ask questions and get in-product help and suggestions based on your actual automation
- Ask for help writing expressions and code based on your needs

To edit with the Assistant:

- Click **Build with Assistant** in the toolbar
- Describe the changes you want to make
- Review and confirm the suggested modifications
 - The chat includes a summary of changes. Click **Reject changes** to revert to the previous version.
- Test the updated automation

Editing on canvas

Make changes directly in the visual designer:

- Select the step or action to modify
- Use the properties panel to adjust settings
- Rearrange steps using drag-and-drop

Note

Every action has a menu of options with common actions like duplicating the action or deleting it.

Best practices for editing:

- Make incremental changes
- Test after each significant modification

Note

Changes are saved automatically to the live version. If you want to deploy those changes, commit and deploy the updated version.

Tip

If you encounter issues while building:

- Use the Automation Assistant to get help
- Check the action documentation for specific requirements

Automation inputs and outputs

Define input and output schemas to create reusable, parameterized automations in Amazon Quick Automate. By defining input and output schemas using the Start and End nodes, you can transform static automations into reusable workflows that accept different data each time they run. Instead of hardcoding values, you define typed input parameters that are provided at runtime and structured output values that are captured when the automation completes. Inputs and outputs work across all invocation methods: manual runs, API calls, and scheduled triggers.

Key benefits of using inputs and outputs:

- **Reusability** – Run the same automation with different data without modifying the automation itself.
- **Type safety** – Amazon Quick Automate validates input data against the schema before execution starts, which prevents invalid data from running.
- **API integration** – Auto-generated schemas enable programmatic discovery and integration with external systems.
- **Observability** – Structured outputs are captured as execution artifacts for audit and review.

Supported data types

The following data types are supported for input and output fields:

- **Text** – Plain text values (e.g., text1, text2)
- **Number** – Numeric values including decimals (e.g., 3.14, 100)
- **Boolean** – True or false values
- **File** – File objects with a maximum size of 5 MB

Start and End nodes

Every automation includes a Start node and an End node, which are blank by default. You can edit these nodes to create a schema, use the inputs in the automation, and update the output data

to return from the automation. The Start node receives input data at runtime, while End nodes (including End Process nodes) collect output data during execution and return it as structured artifacts when the automation completes.

Start

The Start node is the entry point that accepts input parameters when the automation is triggered. It is blank by default.

Properties:

- **Input format** – Defines the schema for input parameters that the automation accepts. Edit using the Input schema editor described in Defining input and output schemas.
- **Input variable** – Variable that stores the input values for the automation as defined in the input format and supplied by the user.

End

The End node is the termination point that collects and returns output values when the automation completes successfully. You can set the output values from variables in the automation using the properties of the End node. Your automation can have multiple End nodes (End process) based on the structure of the automation.

Properties:

- **Output format** – Defines the schema for output parameters that the automation returns. Choose **Edit** and add or modify the schema using the Output schema editor described in Defining input and output schemas.
- **Output data** – Maps values from variables in your automation to the output fields that you defined. For each output variable that you create in the End node (Output schema), specify which automation variable contains the data that you want to return. Use the **Edit output data** editor to configure these mappings in the **Output Value** field.

Defining input and output schemas

Schemas define the structure of data that your automation accepts as input and produces as output. The authoring studio provides a visual form builder for defining input and output schemas directly on the automation canvas.

To define an input schema:

- Open your automation in Amazon Quick Automate.
- Choose the **Start** node on the canvas to open the schema editor. You can also open the schema editor from the properties pane. To do this, choose the **Start** node, and then choose **Edit in Input format** in the properties pane.
- Choose **+Add field** to add an input field.
- For each input field, configure the following properties:
 - **Name** – A unique identifier for the field.
 - **Type** – The type of data the field accepts. For more information, see Supported data types.
 - **Required** – Whether the field must be provided when running the automation.
 - **Default value** (Optional) – This option is active only if the **Required** field is cleared. Default values appear pre-populated in the input form when you run a test or trigger the automation.
 - **Description** – A description of what the field represents.
- Choose **Save** to store the schema with your automation.

To define an output schema:

- Choose an **End** node on the canvas to open the Output schema editor. You can also open the schema editor from the properties pane. To do this, choose the **End** node, and then choose **Edit in Output format** in the properties pane.
- Choose **+Add field** to add an output field.
- For each output field, configure the following properties:
 - **Name** – A unique identifier for the field.
 - **Type** – The type of data for the output.
 - **Description** – A description of what the output represents.
- Choose **Save**.

Note

You can define inputs from the Start node and outputs from any End node. If your automation has multiple End nodes, adding or modifying the schema of one End node modifies all End node schemas.

Using input values from the start node

When your automation has an input schema defined at the Start node, you can access those input values throughout your workflow using the `inputs` dictionary. The runtime automatically validates and populates this dictionary with the values provided when the automation runs.

```
# Access required input fields
value = inputs["field_name"]

# Example usage
customer_id = inputs["customer_id"] # Retrieves a required string input
```

Setting output values in the End node

The End node defines which values from your automation are returned as outputs. You configure this using the **Output Data** property, which provides two interaction modes for setting output values in the editor. Toggle between the modes by choosing the code icon (`</>`) next to the **Output value** field.

Two modes for setting outputs:

- **Variable Selection Mode (Dropdown)** – The default interface displays a dropdown list of all available variables in your automation. Select the variable that you want to assign to each output field.
- **Expression Mode (Code)** – Choose the code icon (`</>`) next to any output value field to switch to expression mode. This mode allows you to enter custom expressions, perform calculations, access nested data, or set literal values.

Defining schemas with the Build with Assistant

Build with Assistant can create or modify input and output schemas directly. Describe your input and output requirements in natural language, and the assistant generates the schema definition for you. Any changes that the assistant makes automatically sync with the Studio visual form builder.

Considerations

Keep the following in mind when you use automation inputs and outputs:

- Input and output schemas are optional. Existing automations without schemas continue to work unchanged.
- File inputs have a maximum size of 5 MB. For larger files, pass the file location (such as an Amazon S3 path) as a text input instead.
- Input and output values are stored with each execution run for observability and audit purposes.
- When you update a schema, you must redeploy the automation for the changes to take effect on deployed runs. Test runs always use the latest draft schema.

You can use the input and output values in the following ways:

- Testing and running the automation from the canvas. For more information, see [Testing automations](#).
- Manually triggering deployed automations. For more information, see [Deploying automations](#).
- Scheduled triggers for deployed automations. For more information, see [Deploying automations](#).

Managing automation versions

Amazon Quick Automate provides version control capabilities to help you track and maintain the history of updates to your automations and easily restore previous versions.

Key concepts

- **Live version** - The current working copy of your automation that you can edit in the canvas. Changes are auto-saved to the live version as you edit.
- **Committed versions** - Read-only snapshots of your automation which can be deployed. Commit a new version of your automation to keep track of significant updates and be able to revert back if needed.
- **Deployed versions** - Committed versions can be deployed and activated to run on a schedule.

Note

Each version maintains its own runtime configuration values. Changes to the Live version's runtime settings don't affect previously committed or deployed versions.

Committing versions

To commit a version:

- Click **Commit** in the canvas toolbar
- Add a descriptive note explaining your changes
- Choose how to increment the version number:
 - Minor version (e.g., 1.05 → 1.06) for smaller updates (default)
 - Major version (e.g., 1.0 → 2.0) for significant changes
- Click **Commit**

Important

You cannot commit a version if the automation has any validation errors. Resolve all errors before attempting to commit.

When you commit a version:

- The committed version becomes a read-only snapshot
- The current runtime configuration is saved with the committed version
- A new Live version is created based on that snapshot to continue editing

Important

Only committed versions can be deployed. The Live version must be committed first before it can be made deployed.

Viewing versions

View versions in two places:

- **Version dropdown in canvas** - Defaulted to the Live version. Choose any previous version from the dropdown to view in read-only mode on the canvas.

- **Versions tab** - Found within Project details. Shows a complete version history including version number, when the version was committed, the user that committed the version, and any version notes. Click on the actions menu to view version details or deploy the version.

Restoring previous versions

To restore a previous version and continue editing it as the Live version:

- Select the version from the version dropdown above canvas
- Choose **Restore live version**
- The selected version replaces your current Live version
- Continue editing the Live version to make any needed adjustments
- Commit as a new version to take a snapshot or deploy the updates

Note

Restoring a version creates a new working copy but doesn't delete any version history.

Best practices

- Add clear version notes to track the purpose of changes
- Commit versions after significant updates
- Test changes before committing
- Review runtime configuration before committing

Automation actions

Actions are the building blocks for your automations. Each action performs a specific task, from reading data to interacting with external systems. This section describes all available actions and how to use them effectively.

Action Categories

- **Agents** - AI-powered actions that use natural language to perform complex tasks like web automation and multi-tool orchestration

- **Quick** - Use native integrations with Amazon Quick features and services directly in your automations
- **Process flow** - Control the flow of your automation with steps, loops, and conditions
- **General** - Common utility actions
- **Web browser** - Interact with websites through clicking, typing, and extracting data
- **Microsoft Excel** - Read, write, and manipulate Excel spreadsheets
- **Data tables** - Work with structured data in memory for processing and analysis
- **Orchestration** - Manage work items through cases and coordinate human-in-the-loop tasks
- **Exception handling** - Define how your automation responds to different types of errors
- **Files** - Manage files used within your automation
- **PDF** - Create, modify, and extract content from PDF documents
- **Code** - Write custom Python code to extend automation capabilities
- **Integrations** - Connect with AWS services and external applications like Salesforce, ServiceNow, and Microsoft 365

Agents

Actions that use AI to understand instructions and perform complex tasks.

- **UI agent** - AI agent for web browser tasks. Used for dynamic and intelligent web automation. Simply write instructions to have it navigate websites, extract data, and generate structured outputs.
- **Custom agent** - AI agent for complex tasks. Create an agent that can understand instructions in natural language and take actions using available tools. Used for tasks that require reasoning, judgement, and dynamic planning.

UI agents

UI agent is a native agent that understands natural language instructions to perform complex browser actions. It can autonomously navigate websites, click, type, read data, and produce structured outputs optimized for downstream automation steps. Example use cases include summarizing products on a webpage or fetching data by navigating websites.

Properties

Title

Name of the step/UI agent

Instructions

In this field you write the prompt for the agent in natural language. Best practices while writing the prompt:

- Be clear and explicit about what you want.
- Structure the prompt. Start with mentioning the 'Task' or 'Role' first and then 'Instructions' to achieve the task with numbered steps
- Add constraints (e.g., only review the products section) and specify when to stop/end (e.g., stop when you find the relevant info)
- Provide positive and negative (don't do this) examples
- Specify length requirements (e.g., less than 100 words) or output format (e.g., date in MM/DD/YY format) clearly

Wrap the text in triple quotes (""") to write multiline prompts. For example:

```
"""Task: Locate the company's latest annual report.  
* Visit the provided URL.  
* Look for the annual report. The report may be titled 'Annual Report', 'Financial  
Report', 'Year in Review', or similar variations..."""
```

Structured Output (optional)

Agent Response: Name of the variable to assign the output of this operation

How to configure structured output fields

Adding fields

- Click Add field to create a new output field
- Enter the Output name - this becomes the JSON property name
- Select the Type from the dropdown
- Check Required if the field must always be present
- Add a Description to guide the AI agent

Field types

- *String* - Text values (Names, descriptions, summaries)
- *Number* - Numeric values (Counts, scores, percentages)
- *Boolean* - True/false values (Status flags, yes/no questions)
- *Object* - Nested structure (Complex data groupings)
- *Array* - List of items (Tags, categories, multiple values)
- *File* - File references (Document attachments, images)
- *Data table* - Tabular data (Structured datasets, reports)

Working with complex types

Objects and Arrays can contain nested fields:

- Click the expand arrow (►) next to Object or Array fields
- Use Add field within the nested structure
- Keep nesting to 2-3 levels maximum for optimal performance

Example configuration

Here's a simple configuration for summarizing customer feedback:

```
{
  "orderId": "12345",
  "numberOfOrders": 3,
  "hasShipped": true,
  "orderDetails": {
    "quantity": 2,
    "productName": "ABC",
  },
  "tags": ["electronics", "urgent"]
}
```

This structure would be configured as:

- orderId (String, required)
- numberOfOrders (Number, required)
- hasShipped (Boolean, required)

- **orderDetails** (Object, required)
 - **quantity** (Number, required)
 - **productName** (String, required)
- **tags** (Array of Strings, Optional)

Best practices

- Use descriptive field names - Help the AI understand what data to extract
- Add clear descriptions - Provide context for complex fields
- Mark critical fields as required - Ensure essential data is always present
- Limit nesting depth - Keep structures simple for better performance
- Test your configuration - Verify the output matches your expectations by running the agent step and verifying the response.

Important notes

- **JSON Knowledge:** Unfamiliar with JSON? Learn the basics at json.org
- **No validation:** Currently, the system doesn't validate output structure - ensure your automation handles missing or malformed data

Custom agents

Custom agent is an intelligent action that processes natural language inputs to automate complex steps. It uses integrated tool-calling capabilities with integrations as its primary tool interface. It also offers extensibility to use Code as a tool and other native actions like human-in-the-loop task. The agent delivers structured, predictable outputs optimized for downstream automation steps.

Properties

- **Title:** Name of the step/custom agent
- **Mode:** A mode defines how the agent operates based on your use case. The three available modes are Fast, Pro, and Custom. Fast is best for simple tasks like summarization, classification, and high-volume automations. Pro is ideal for complex tasks that involve reasoning and orchestration of multiple tools or actions. Fast and Pro are fully managed modes that require no pre-setup in advance. Custom Mode requires a Bedrock runtime connector and lets you select

the model you want to use. For more information, see [Using Custom Models in Custom Agent \(Bring your own bedrock model\)](#). Custom Mode is ideal when you have a prompt fine-tuned for a particular Bedrock model, need a specific Bedrock model for the Agent, or want to use your own custom or fine-tuned model hosted on Bedrock. In Custom Mode, you bring your own model from Bedrock through an integration. Model inference is billed separately to the account associated with that Bedrock integration.

- **Instructions:** In this field you write the prompt for the agent in natural language. Best practices while writing the prompt:
 - Be clear and explicit about what you want.
 - Structure the prompt. Start with mentioning the 'Task' or 'Role' first and then 'Instructions' to achieve the task with numbered steps
 - To improve tool-call accuracy and guide the Agent, clearly specify in the prompt which tool to use at each step, if applicable.
 - Specify length requirements (e.g., less than 100 words) or output format (e.g., date in MM/DD/YY format) clearly
 - Wrap the text in triple quotes (""") to write multiline prompts. For example:

```
"""You are content summarization agent.  
Summarize the last two paragraphs of the provided text, focusing only on the main  
conclusion."""
```

- **Actions:** Action is a tool that enables the AI agent to interact with external systems or perform specific tasks. This is optional. You can run the custom agent without any actions. The following actions can be used in the custom agent.
 - **General Actions**
 - **Create user task** - If enabled, this tool allows the Agent to trigger a Human-in-the-Loop (HITL) task when it gets stuck during execution. The Agent pauses and waits for human input. The HITL task is visible in the task center. For best results, specify in the prompt exactly when the Agent should invoke HITL. This action is selected by default. The automation runs until the task is finished.
 - **Code** - The Code action generates and executes Python code within a restricted Python environment, same as code actions. It solves tasks involving calculations, data manipulation, and file processing. Unlike code generators, it actively creates and runs scripts to accomplish objectives. It works with Excel, PDF files, various data formats, and available integrations.
 - **Key Capabilities:**

- **File Operations:** Process multi-tab Excel files, extract content, perform date calculations, apply conditional formatting, and upload results to S3
- **Data Transformation:** Convert between JSON and table formats, transpose data, rename columns, and join tables
- **Advanced Computations:** Generate numerical sequences and perform automated validation
- **Integrations:** If you have added specific integration actions — such as Salesforce, MS Exchange, or Bedrock—to your automation group, their corresponding actions appear here to be use in the custom agent. The author can then select the relevant actions to use as tools for the agent.

List of integrations which can be used as tools/actions in the custom agent

- Amazon S3
- Amazon Bedrock Data automation
- Amazon Comprehend
- Amazon Textract
- Custom REST API
- Custom MCP connector
- Microsoft Outlook
- Salesforce
- **Structured Output (optional)**

Configure your AI agent to return structured JSON output that downstream steps can process. This feature is ideal for text summarization, report generation, data transformation, and extracting statistics from unstructured content. This is an optional field. If you do not define structured output, the agent returns output in natural language by default. Use structured output when your output has a defined structure, such as a list, data table, or JSON.

 Note

The structured output configuration for Custom agents follows the same format as UI agents. Refer to the UI agent structured output section for detailed configuration instructions.

- **Agent response:** Name of the variable to assign the output of the agent. The response follows your structured output format in a JSON schema if defined, otherwise is a free-form text.

Using Custom Models in Custom Agent (Bring your own bedrock model)

Integrate your desired or custom fine-tuned models hosted in AWS Bedrock with Quick Suite automation workflows.

Before you begin, ensure you have the following:

- A fine-tuned model deployed and accessible in AWS Bedrock
- Quick Suite Admin access for creating connectors
- An IAM role with Bedrock permissions for invoking models
- Your model ID (for example, `us.anthropic.claude-3-5-sonnet-20241022-v2:0`)

Step 1: Create a Bedrock Runtime Action integration by following the detailed instructions in [AWS service connectors](#)

Step 2: Set Up Your Automation Group

Create an automation group and connect the integration:

- **Create an automation group** - Follow the detailed instructions in [Setup tasks](#)
- **Configure integrations** - Follow the detailed instructions in [Setup tasks](#)
- Once configured, the connector appears in your available assets list

Step 3: Configure a Custom Agent

Add and configure a custom agent to use your fine-tuned model:

- Within your automation workflow, add a custom agent
- Configure the following agent settings:
 - **Agent Title:** Enter a descriptive name for your agent
 - **Instructions:** Enter custom prompts tailored to your use case
 - **Mode:** Select Custom
 - **Connector:** Choose your Bedrock Runtime connector (required when Custom mode is selected)

- **Custom Model:** Enter your model ID (for example, `us.anthropic.claude-3-5-sonnet-20241022-v2:0`) - required when Custom mode is selected

Next Steps

Once configured, your custom agent uses the fine-tuned model to process requests according to the instructions you provided. You can now incorporate this agent into your Quick Automate workflows.

Note

Ensure your model ID is correctly formatted and matches the model deployed in your AWS Bedrock account. You can find your model ID in the AWS Bedrock console under your provisioned models.

Custom agent testing

Custom agent testing enables you to test individual agents independently from the complete automation workflow. This capability helps you validate agent behavior, debug prompts, and iterate more efficiently without executing the entire workflow.

Prerequisites

- An automation workflow with at least one configured custom agent
- Appropriate permissions to run automations in your workspace

Start a test

- In the workflow canvas, hover over the agent card you want to test
- Choose the **Unit test** button that appears at the top of the card
- In the variable collection window that opens, review the automatically detected variables from your agent's prompt
 - The prompt preview displays all detected variables with highlighting
- Enter a value for each variable
 - Values must use valid expression syntax
 - If a value contains invalid syntax, an error message appears and prevents test execution

Monitor test execution

During test execution, you can monitor progress in the audit panel. The test skips all preceding workflow steps and executes only the selected agent. You get the same logging experience as a full workflow run.

Review test results

After the test completes, review the following information in the Test panel:

- Metrics Card (Monitor Tab in the Test panel)
 - Total execution time
 - Number of tools used
 - Number of tasks created
- Logs in between
- Watch Variables Tab (Variables accordion in the Test panel)
 - Input - View input variables and their values
 - Output - Examine output results from the agent execution
 - For structured outputs, choose **View Details** to open the View Output dialog box:
 - Fields Tab - Navigate data using the tree structure view
 - Fields - Highlight corresponding values by selecting tree nodes in Fields tab
 - Output fields - Corresponding values for the JSON keys

Using Custom agent with Build with Assistant

The current tenet for custom agent is it has to be specifically mentioned to consistently get it invoked, here are the things needed in the prompt to make it appear:

```
- Function names: `use_inline_agent`  
- Representation names: "Custom Agent", "Inline Agent" # use `use_inline_agent`  
- Generic terms: "agentic skills" # default to `use_inline_agent`
```

Otherwise, the model is preferred to author the workflow deterministically.

Although, in practice, when no appropriate actions are available, planner might pick custom agent as a workaround. But to consistently invoke custom agent in the workflow, the above phrases are encouraged to use in the prompt.

Examples of agent use cases

Use Case 1: Email Classification and Assignment Agent

Role: You are an Email Categorization and Assignment Agent

Instructions: Follow these steps:

- Step 1: Classify the incoming email based on the Category column of the provided reference table as knowledge
- Step 2: Use the email system to send a notification:
 - From: [system_email]
 - To: [team_distribution_email]
 - Subject: [Classification Result]
 - Body: Include a brief summary explaining the classification reasoning and key points from the original email
- Step 3: For all valid categories (except 'unknown'), create a new case in Salesforce with:
 - Subject: [Original Email Subject]
 - Description: Summarized issue from email body
 - Priority: Based on content urgency (High/Medium/Low)
 - Type: Select appropriate type (Question/Problem/Feature Request/Other)
 - Status: 'New'
 - Category: [Classification result from Step 1]
- Step 4: If classified as 'unknown':
 - Escalate to supervisor for manual review
 - Add note explaining why classification was uncertain
 - Based on the category received from the supervisor, follow step 2 and 3 and stop
 - If the category received from the supervisor is unknown or invalid, stop

Using knowledge bases with custom agents

In Amazon Quick Automate, you can connect knowledge bases to custom agents to enable AI-powered retrieval and question answering over your organization's documents. By linking a Quick space to your automation group, custom agents can search and retrieve information from the knowledge bases within that space.

Use this for automations that need to reference organizational knowledge — such as answering questions from policy documents, summarizing reports, or classifying content based on reference data.

Knowledge bases index your documents for semantic search, so the custom agent retrieves only the most relevant passages rather than processing entire files. This makes retrieval faster and more accurate, especially across large document sets.

Prerequisites

- A Quick space with one or more knowledge bases configured. For more information, see [Organize, collaborate, and share resources with spaces in Amazon Quick](#).
- An automation group where you are an owner
- Owner-level access to the space you want to link

Link a space to your automation group

Before a custom agent can access knowledge bases, you must link the space that contains those knowledge bases to your automation group. Linking a space grants the automation group permission to access the knowledge bases and files within that space.

To link a space to an automation group:

1. In the **Automations** tab, go to the **Projects** page.
2. Choose **Groups** and select the group you want to attach the space to.



Tip

You can also choose **Create group** to create a new automation group.

3. In the **Assets** section, choose **Add**, and then choose **Spaces**.
4. Select the space that contains the knowledge bases you want to use, and then choose **Add**.

The space now appears in the automation group's connections list. Custom agents in this automation group can access the knowledge bases and files within the linked space. Other resources in the space are not available to automations.

Note

If the space shows **Limited access** after being added, it means not all knowledge bases are shared with the automation group. This can happen if knowledge bases were added to the space after linking, or if not all knowledge bases were shared initially. To resolve this, refresh the space connection to share all resources with the automation group. You do not need to reconfigure the Knowledge tab on individual custom agents.

Add knowledge to a custom agent

After linking a space to your automation group, you can configure a custom agent to use the knowledge bases within that space. The workflow must be in the same automation group where you attached the space.

To add knowledge to a custom agent:

1. In the workflow builder, add a **Custom Agent** step. You can either drag and drop a Custom Agent node onto the canvas, or chat with the automation assistant to build this step.
2. In the agent properties panel, choose **Knowledge**, and then choose **Add**.
3. A picker opens showing available spaces linked to the automation group. Select one or more spaces that contain the knowledge bases you want the agent to use.
4. Choose **Save**.

The custom agent can now search and retrieve content from the knowledge bases in the selected spaces when the automation runs.

When you attach a space, all knowledge bases inside that space automatically become available to the agent. You don't need to attach each knowledge base individually. At runtime, the agent queries each knowledge base independently and combines the results in its response.

Note

If the automation group owner loses access to a specific knowledge base within the space, that knowledge base is skipped during queries and the workflow editor displays a warning badge on the space attachment.

Writing instructions for knowledge base queries

When a custom agent has knowledge bases attached, it automatically searches and retrieves relevant content based on your instructions. Write instructions that clearly describe what information the agent should find or how it should use the knowledge base content.

Best practices:

- Be specific about what information to retrieve or summarize
- Reference the type of content you expect the agent to find (for example, "Search the policy documents for..." or "Find information about...")
- Specify how the agent should use the retrieved information in its response
- Include fallback instructions for when the knowledge base doesn't contain relevant content

Example: Customer inquiry agent with knowledge base

The following example shows how to configure a custom agent that uses a knowledge base to answer customer inquiries based on company documentation.

Setup:

- A space containing a knowledge base with product documentation and FAQ content
- The space is linked to the automation group
- The space is added as knowledge to the custom agent

Instructions:

```
""""You are a customer support agent.
```

```
Task: Answer the customer inquiry using information from the knowledge base.
```

```
Instructions:
```

1. Search the knowledge base for information relevant to the customer's question.
2. Provide a clear, concise answer based on the retrieved content.
3. If the knowledge base does not contain relevant information, respond with:
 "I don't have enough information to answer this question. Please escalate to a human agent."

Constraints:

- Only use information found in the knowledge base. Do not make up answers.
- Keep responses under 200 words.
- Include the source document name when referencing specific information."

Structured output:

```
{
  "answer": "The response to the customer inquiry",
  "sourceDocument": "Name of the document used",
  "confidence": "high/medium/low",
  "escalationNeeded": false
}
```

Using datasets with custom agents

In Amazon Quick Automate, you can connect datasets to custom agents so that an agent can answer natural-language questions about your data. The underlying query returns a list of data tables and a string output from the API call. You can then convert this output into a desired data structure, either by prompting the custom agent or by using the structured output feature of the custom agent.

Dataset Q&A (querying datasets) is an agentic action that you use with the custom agent. A custom agent can query multiple datasets, but it runs one query at a time.

Use datasets for data-driven question-answering automations. For example, you can build an agent that answers questions such as "which customers have unpaid invoices over \$10,000" by querying an invoices dataset and returning the matching rows.

Prerequisites

- Owner-level access to the dataset you want to attach
- A dataset that does not use row-level security (RLS) or column-level security (CLS). Datasets with RLS or CLS are not supported.
- An automation group where you are an owner
- Datasets attached to the automation group for run-as-service mode, or datasets shared with your user ID for run-as-user mode

Attach a dataset to your automation group

Before a custom agent can query a dataset, you must attach the dataset to your automation group. You must already have access to the dataset as an owner. Attaching the dataset to the automation group does not grant you access to it.

To attach a dataset to an automation group:

1. In the **Automations** tab, go to the **Projects** page.
2. Choose **Groups** and select the group you want to attach the dataset to.

Tip

You can also choose **Create group** to create a new automation group.

3. In the **Assets** section, choose **Add**, and then choose **Datasets**. The picker opens with **Actions**, **Credentials**, **Spaces**, and **Datasets** tabs.
4. On the **Datasets** tab, select the dataset you want to use, and then choose **Add**. You can search your Quick assets. Datasets that you cannot use appear disabled, and you can hover over one to see the reason, such as needing to be an owner.

The dataset now appears in the automation group's assets and connections list. Custom agents in this automation group can query it.

Note

You must be the owner of the dataset to use it. Datasets that use row-level security (RLS) or column-level security (CLS) are not supported.

Add datasets to a custom agent

After you attach one or more datasets to your automation group, you can configure a custom agent to query them. The workflow must be in the same automation group where you attached the datasets. After you add the Custom Agent step, you add the dataset from the agent's properties panel on the **Datasets** tab.

To add datasets to a custom agent:

1. In the workflow builder, add a **Custom Agent** step. You can either chat with the automation assistant to build this step, or drag and drop a Custom Agent node onto the canvas.
2. In the agent properties panel, choose the **Datasets** tab.
3. Add one or more datasets directly on the **Datasets** tab.
4. Choose **Save**.

A custom agent can use multiple datasets, but it runs one query at a time. If you connect a Quick space that contains a dataset, the agent can query that dataset as well. For more information about spaces, see [Organize, collaborate, and share resources with spaces in Amazon Quick](#).

Query a dataset with the agent

After you add a dataset to the custom agent, you can ask the agent a question about the data. When the automation runs, or when you test the agent, the agent queries the attached dataset and returns a response based on the results.

To query a dataset:

1. In the agent **Instructions**, write a prompt that asks a question about the data. For example, ask a natural-language question about records in the dataset.
2. Run the automation, or use the agent unit test to run the agent. For more information about testing an agent, see [Custom agent testing](#).

The agent queries the dataset and returns the answer.

To return the results in a specific format that downstream steps can process, such as a JSON object or a data table, configure the Structured Output option for the agent. Structured output is optional and returns structured JSON. The structured output configuration for custom agents follows the same format as UI agents. In your prompt, instruct the agent to produce output that matches the structured output schema.

Use datasets in a space with a custom agent

Instead of attaching datasets directly on the agent's **Datasets** tab, you can add your datasets to a Quick space, attach the space to your automation group, and then attach the space to the custom agent on the **Knowledge** tab. The agent can then query the datasets in that space based on your instructions. This approach is useful when you already organize datasets, and other knowledge, in a space.

To use datasets in a space:

1. Add your dataset or datasets to a Quick space. For more information about spaces, see [Organize, collaborate, and share resources with spaces in Amazon Quick](#).
2. Attach the space to your automation group. In the **Automations** tab, on the **Projects** page, choose **Groups** and select your group. In the **Assets** section, choose **Add**, and then choose **Spaces**. Select the space and choose **Add**.
3. In the workflow builder, add a **Custom Agent** step. You can either chat with the automation assistant to build this step, or drag and drop a Custom Agent node onto the canvas.
4. In the agent properties panel, choose the **Knowledge** tab, choose **Add**, select the space that contains your datasets, and then choose **Save**.
5. In the agent **Instructions**, write a prompt that asks a question about the data in the space's datasets.

When the automation runs, the agent queries the datasets in the attached space and returns the results. The datasets must still meet the dataset requirements. You must be an owner of each dataset, and datasets that use row-level security (RLS) or column-level security (CLS) are not supported.

Writing instructions for dataset queries

When a custom agent has datasets attached, it automatically determines the relevant dataset and constructs the query based on your instructions. Write instructions that clearly describe what data the agent should retrieve and how it should present the results.

Best practices:

- Be specific about what data to retrieve, such as the columns, records, or aggregations you expect in the results.
- Include filters, calculation definitions, or metric logic in the prompt to improve query accuracy.
- Specify the output format so that downstream steps can process the results.

Example: Invoice analysis agent with a dataset

The following example shows how to configure a custom agent that queries a dataset to answer questions about outstanding invoices.

Setup:

- A dataset that contains customer invoice records, including customer name, invoice amount, and payment status
- The dataset is attached to the automation group
- The dataset is added to the custom agent on the Datasets tab

Instructions:

```
""""You are an invoice analysis agent.
```

```
Task: Answer the user's question by querying the invoices dataset.
```

```
Instructions:
```

1. Identify the relevant dataset and construct a query for the user's request.
2. For the request "Show me customers with unpaid invoices over \$10,000", filter for records where payment status is unpaid and invoice amount is greater than 10000.
3. Return the results as a data table with customer name and invoice amount.

```
Constraints:
```

- Only use data found in the dataset. Do not make up values.
- Sort the results by invoice amount in descending order."""

Output:

The agent returns a data table similar to the following:

Customer name	Invoice amount
Example Corp	24,500
AnyCompany Retail	15,200
AnyCompany Financial	11,750

Quick

With Quick actions, you can use native integrations with Amazon Quick features and services directly in your automations. These actions work within the Quick environment without requiring external authentication or complex infrastructure setup. You can interact with other Quick capabilities while using unified governance, permissions, and data management.

Shared File Actions

Shared file storage provides a centralized location for managing files used by your automations. Files uploaded by one automation are immediately available to other automations in the same automation group. This eliminates the need for external storage solutions for lightweight file management scenarios.

Key Benefits:

- Share files across multiple automations within the same automation group
- Avoid external storage setup (for example, Amazon Simple Storage Service or SharePoint)
- Upload files manually by using the built-in file management interface (Select File)

The following two actions are available for shared file operations.

- **Download shared file** – Download a shared file for use in the automation.
- **Save as shared file** – Upload a file from the automation to shared storage.

Download Shared File

Retrieves a file from shared storage for use in your automation. Use this action to retrieve files that were previously uploaded either manually or by another automation in the same automation group.

Properties:

- **File to download** (Required): The file to retrieve from shared storage. The File to download property includes a folder icon (#) next to the field. Choose this icon to open the Select File editor, which provides an interface to:
 - *View all files*: See all files currently stored in your automation group's shared storage. Files appear as a list with Name, Last modified by, and Last modified time.
 - *Actions*: Each file has an Actions menu (:) with the following options:
 - *Download*: Save the file to your local machine
 - *Delete*: Remove the file from shared storage
 - *Upload files*: Choose **Upload** to manually add new files to shared storage.
 - *Select and Save*: Choose a file from the list and choose Save to use it in your action, or choose Cancel to close without selecting.

- **Output Variable** (Required): Variable name to store the downloaded file object.

Save as shared file

Uploads a file to shared storage, making it available to other automations in the automation group. Use this action to save output files generated by your automation. Other automations can then access these files, or you can download them manually.

Properties:

- **File to save** (Required): The file object to upload to shared storage (for example, an Excel file, PDF, or other file type generated by your automation).
- **Save as filename** (Optional): The name to use when saving the file in shared storage. If not provided, the action uses the original filename.
- **Replace existing file** (Optional): When checked, the existing file with the same name is replaced by the file being uploaded

Access the files manually

You can manually access and download files in Amazon Quick Automate using either of the following methods.

From the authoring studio (Action)

- Choose the folder icon (#) next to the **File to download** property for the **Download shared file** action to open the **Select File** dialog box. This displays the list of files available for the automation group.
- Choose the three dots (:) under **Actions**, and then choose **Download** to save the file to your local machine.

From the automation group

- Open your automation group in Quick Automate.
- Choose the **Shared files** tab to view the list of available files.
- Choose the three dots (:) under **Actions**, and then choose **Download** to save the file to your local machine.

Limitations

Shared file storage has the following limits per automation group.

Limit	Value	Scope
Maximum number of files	250	Per automation group
Maximum total storage	1 GB	Per automation group
Maximum individual file size	200 MB	Per file

Note

Supported file types: HTML, TXT, JSON, CSV, PDF, XLSX, XLS, JPEG/JPG, PNG, DOCX, DOC, TIFF/TIF, and ZIP.

Best Practices

- **Use descriptive filenames:** Include dates or identifiers in filenames to avoid confusion (for example, "report_2026-03-31.xlsx").
- **Clean up old files:** Regularly delete files that are no longer needed by using the Select File editor.
- **Consider file size limits:** Individual files can be up to 200 MB, and total shared storage is limited to 1 GB per automation group. For large-scale file storage requirements, consider using Amazon S3 integration instead.
- **Use consistent naming conventions:** Establish naming patterns across your automation group for easier file management.

Notes

Keep the following considerations in mind when using shared file actions.

- Quick Automate stores files at the automation group level. This means:
 - All automations within the same automation group can access the same shared files.

- Files uploaded by one automation are immediately available to other automations in the group.
- The file count limit (250 files) and total storage limit (1 GB) apply at the automation group level. All automations in the group share these limits.

Examples

Example 1: Processing a shared configuration file

Scenario: Multiple automations need to read the same configuration file.

1. Download Shared File
 - File to download: "config.xlsx"
 - Output Variable: configFile
2. Read Excel Range
 - File: configFile
 - Sheet name: "Settings"
 - Output Variable: configData
3. Continue with automation logic using configData

Example 2: Sequential processing across automations

Scenario: Automation A processes data and saves it. Automation B retrieves and continues processing.

Automation A:

1. [Process data and create output file]
2. Save as Shared File
 - File to save: processedData
 - Save as filename: "processed_orders.csv"

Automation B (runs after Automation A):

1. Download Shared File
 - File to download: "processed_orders.csv"
 - Output Variable: ordersFile

2. Continue processing the orders data

Example 3: Manual upload and automation processing

Scenario: A user manually uploads a file that an automation will process.

Manual steps:

- The user opens the Select File editor from any Download Shared File action.
- The user chooses Upload and selects "customer_list.xlsx" from their computer.
- The user chooses Save to confirm the upload.

Automation steps:

1. Download Shared File
 - File to download: "customer_list.xlsx"
 - Output Variable: customerFile
2. Process customer data

Process flow

- **Process step** - Groups logically related actions. Used for organizing your process.
- **If decision** - Checks a true/false condition. The process is split into two branches based on the result.
- **Loop through items** - Repeats actions for each item. Used for processing multiple items like emails, files, or data rows one at a time.
- **Loop while true** - Repeats while a condition is met. The condition is checked before each loop and must be True for the loop to continue.
- **End process** - Ends the current run successfully. Optionally, provide a custom message or value as the final outcome of the process.

General

- **Save value** - Assigns a value to a variable. Used to save information that you can reference later in your automation.

- **Write log message** - Adds a message to the logs. Useful for tracking progress, troubleshooting issues, or recording important events.
- **Wait** - Pauses the automation. After the specified amount of time, the automation will continue with the next action.

Web browser

Overview

Browser actions provide a comprehensive automation toolkit for performing web interactions. Browser sessions are automatically managed — opened, maintained, and closed as needed. Quick Automate supports flexible element identification through AI-powered natural language descriptions and precise XPath selectors. Key capabilities include intelligent element targeting with fallback mechanisms, self-healing features that automatically handle popups and retry failed actions, data extraction including structured table content, and a full suite of interaction methods from basic clicks and typing to advanced AI-driven multi-step tasks.

Element identification methods

- **Description method:** Uses AI to identify elements using natural language descriptions (e.g., "Log In button"), making automations more resilient to UI changes.
- **XPath selector method:** Uses precise XPath expressions for exact element targeting (e.g., `//button[text()='Log in']`).
- **Combined approach:** When both methods are enabled, XPath is tried first before falling back to the description method, combining precision with adaptability.

Self-Healing features available for type-ins and clicks

- **HandlePopups:** Automatically detects and handles unexpected dialogs, cookie notices, and overlays, preventing blocking elements from interrupting workflows.
- **ValidateAndRetry:** Verifies action success through screenshot comparison and retries if needed, reducing false failures.

Browser Session Management

Browser instances are automatically created, maintained throughout workflows, and properly cleaned up when complete. The `restart_browser` action allows fresh sessions when needed for scenarios such as login timeouts or session expiration.

Start Browser Session

Opens a new browser session. Actions inside is scope interact with this browser. The browser closes automatically when finished.

Properties:

- **Action Title** (text-only): The title of the action displayed in the process visualization (default: Browser)

Note

1. All browser actions must be placed within a browser session scope
2. Each browser scope is a browser session. Browser automatically closes when actions inside a scope completes.

Go to Webpage

Navigates to a specific URL. Used to go to a new webpage in the current browser tab.

Properties:

- **URL** (required): The web address to navigate to (e.g., "example.com") Note: Please provide the complete URL when working with this.

Click

Clicks on a webpage element. Used to interact with buttons, links, or other clickable elements.

Properties:

- **Find Element with Description** (radio button): Use natural language description to identify the element with AI (default: ON)

- **Description Field** (conditionally required): Natural language description of the browser element to click (e.g., "Submit button"). Don't use the action in the element, just the description of the element
- **Find Element with Selector** (radio button): Use XPath expression to identify the element (default: OFF)
 - **XPath Field** (conditionally required): XPath expression (e.g., "//input[@id='xyz']")
- **Self-healing** (checkbox): Use AI to handle unexpected popups automatically (default: OFF, optional)
- **Mouse Button** (dropdown): Which mouse button to click - Left, Middle, or Right (default: Left, optional)
- **Click Type** (dropdown): Single or Double click (default: Single, optional)
- **Get File Download** (radio button): Enable if click downloads a file (default: OFF, optional)
- **Downloaded File** (output, conditionally required): Variable name storing the downloaded file (default: downloaded_file)

Element Identification:

- At least one element identification method must be enabled
- When both methods are enabled, XPath selector is tried first before falling back to description

Enter Text

Types text into an input field. Used to fill in forms, search boxes, or other text inputs on webpages.

Properties:

- **Text to Enter** (required): The text you want to type into the field (e.g., "Order #12345")
- **Find Element with Description** (radio button): Use natural language description to identify the element with AI (default: ON)
 - **Description Field** (conditionally required): Natural language description (e.g., "Order number field")
- **Find Element with Selector** (radio button): Use XPath expression to identify the element (default: OFF)
 - **XPath Field** (conditionally required): XPath expression (e.g., "//input[@id='xyz']")

- **Self-healing** (checkbox): Use AI to handle unexpected popups automatically (default: OFF, optional)
- **Replace Existing Text** (checkbox): Clear any pre-existing text in the field before typing (default: ON)

Element Identification:

- At least one element identification method must be enabled. Recommended to use element description and move to element selector only if required.

Enter Credential

Types a username or password. Used to securely sign in to a website using saved credentials.

Properties:

- **Credential** (dropdown, required): Choose which saved credential to use from automation group
- **Value to Enter** (dropdown, required): Choose whether to enter Username or Password (default: Username)
- **Find Element with Description** (radio button): Use natural language description to identify the element with AI (default: ON)
 - **Description Field** (conditionally required): Natural language description (e.g., "Username field")
- **Find Element with Selector** (radio button): Use XPath expression to identify the element (default: OFF)
 - **XPath Field** (conditionally required): XPath expression (e.g., "//input[@id='username']")
- **Self-healing** (checkbox): Use AI to handle unexpected popups automatically (default: OFF, optional)

Element Identification:

- At least one element identification method must be enabled

Select from Dropdown

Chooses a value from a dropdown. Used to make selections in forms, filters, or other dropdown menus on webpages.

Properties:

- **Value to Select** (required): The option you want to select from the dropdown field (e.g., "Complete")
- **Find Element with Description** (radio button): Use natural language description to identify the element with AI (default: ON)
 - **Description Field** (conditionally required): Natural language description (e.g., "Order status dropdown")
- **Find Element with Selector** (radio button): Use XPath expression to identify the element (default: OFF)
 - **XPath Field** (conditionally required): XPath expression (e.g., "//select[@id='xyz']")
- **Self-healing** (checkbox): Use AI to handle unexpected popups automatically (default: OFF, optional)

Element Identification:

- At least one element identification method must be enabled

Extract Text

Reads a value from a webpage. Used to capture the text from an individual field.

Properties:

- **Find Element with Description** (radio button): Use natural language description to identify the element with AI (default: ON)
 - **Description Field** (conditionally required): Natural language description (e.g., "Total amount field")
- **Find Element with Selector** (radio button): Use XPath expression to identify the element (default: OFF)
 - **XPath Field** (conditionally required): XPath expression (e.g., "//input[@id='xyz']")
- **Self-healing** (checkbox): Use AI to handle unexpected popups automatically (default: OFF, optional)
- **Extracted Text** (output): Variable name storing the extracted text (default: extracted_text)

Element Identification:

- At least one element identification method must be enabled

Extract Table

Reads data from a webpage table. Used to capture structured information found in tables.

Intelligently extracts structured data from web pages using natural language targeting or XPath selectors. Features include AI-enhanced extraction, multiple attribute extraction (text, links, etc.), automatic column generation for attributes, and data table output for programmatic manipulation.

Properties:

- **Find Element with Description** (radio button): Use natural language description to identify the element with AI (default: ON)
 - **Description Field** (conditionally required): Natural language description (e.g., "Order details table")
- **Find Element with Selector** (radio button): Use XPath expression to identify the element (default: OFF)
 - **XPath Field** (conditionally required): XPath expression (e.g., "//table[@id='xyz']")
- **Self-healing** (checkbox): Use AI to handle unexpected popups automatically (default: OFF, optional)
- **Extract Multiple Pages** (radio button): Continue extracting data from additional pages if they exist (default: OFF, optional)
- **Next Page Button Selector** (conditionally displayed): XPath expression for next page button (e.g., "//*[@id='nextButton']", optional)
- **Intelligent Data Extraction** (checkbox): Use AI to extract and structure data from complex tables (default: ON). Do not use for large table sizes of 100 rows or more. This option should be used for complex tables or table-like structures, where normal table extraction from a web browser fails
- **Data to Extract** (multi-select dropdown): Choose what to extract - Text, Links (Href), or Element selectors (XPath) (default: Text, optional)
- **Extracted Table** (output): Variable name storing the extracted data table (default: extracted_table)

Element Identification:

- At least one element identification method must be enabled

Usage Notes:

- Additional data types are saved in separate columns (e.g., columnName_Href for Links)
- Disable intelligent extraction for tables with 100+ rows for better performance

Take Screenshot

Captures an image of the webpage. The screenshot is saved to a file to be used later in your process.

Properties:

- **File Name** (optional): Name for the saved image file. Auto-generated if empty (e.g., "Screenshot image")
- **File Type** (dropdown): Choose the image format - PNG or JPEG (default: PNG, optional)
- **Screenshot File** (output): Variable name storing the screenshot file (default: screenshot_file)

Visual Q&A

Answer questions about a webpage. Uses AI to visually analyze the page and answer true/false questions about the content.

Properties:

- **Question** (required): Enter your true/false question about the webpage content (e.g., "Is the product in-stock?")
- **Self-healing** (checkbox): Use AI to handle unexpected popups automatically (default: OFF, optional)
- **Answer** (output): Variable name storing the true/false response (default: is_true)

Click and Upload Files

Uploads files to a webpage. Used to automate file pickers by clicking on a file upload button and choosing the files to upload.

Properties:

- **Files to Upload** (required): The list of files to be uploaded, typically stored in a variable (e.g., [my_file1, my_file2])
- **Find Element with Description** (radio button): Use natural language description to identify the element with AI (default: ON)
 - **Description Field** (conditionally required): Natural language description (e.g., "Upload button")
- **Find Element with Selector** (radio button): Use XPath expression to identify the element (default: OFF)
 - **XPath Field** (conditionally required): XPath expression (e.g., "//input[@id='xyz']")
- **Self-healing** (checkbox): Use AI to handle unexpected popups automatically (default: OFF, optional)
- **Mouse Button** (dropdown): Which mouse button to click - Left, Middle, or Right (default: Left, optional)
- **Click Type** (dropdown): Single or Double click (default: Single, optional)

File Requirements:

- Must be an array of media/file objects

Element Identification:

- At least one element identification method must be enabled

Enter Keystroke

Sends a key command or shortcut. Used to automate typing individual keys or key combinations on a webpage.

Properties:

- **Special Keys** (dropdown, required): Select modifiers like Ctrl, Alt, Shift or singular keystrokes like Enter, Tab, Pagedown
- **Command Key** (optional): Enter the key to be used in combination with the Special key (e.g., "c" for Ctrl+c)

- **Find Element with Description** (radio button): Use natural language description to identify the element with AI (default: ON)
 - **Description Field** (conditionally required): Natural language description (e.g., "Search bar")
- **Find Element with Selector** (radio button): Use XPath expression to identify the element (default: OFF)
 - **XPath Field** (conditionally required): XPath expression (e.g., "//input[@id='xyz']")
- **Self-healing** (checkbox): Use AI to handle unexpected popups automatically (default: OFF, optional)

Element Identification:

- Both element toggles can be OFF - keystroke is sent to the current active field if no element is defined

Wait for Element

Checks for an element to appear. Used to wait for changes or validate conditions on a webpage. Returns true if the element is found within the maximum wait time, false if not.

Properties:

- **Action Title** (text-only): Display name for the action
- **Element Selector** (required): XPath expression to define the element you want to check (e.g., "//input[@id='xyz']")
- **Max Wait Time** (optional): Maximum time to wait in seconds for the element to appear (default: 30)
- **Element Found** (output): Variable name storing whether the element was found (True/False) (default: is_found)

Validation:

- Max wait time must be an integer or decimal

Find Matching Elements

Gets elements matching a pattern. Used to get a list of similar elements on a webpage that you can process one by one.

Properties:

- **Element Selector** (required): XPath expression to define the pattern for elements you want to get (e.g., "//div[@class='abc']")
- **List of Elements** (output): Variable name storing the list of found elements (default: `element_list`)

Usage Notes:

- The list is empty if no matching elements are found
- Use for processing multiple similar elements like search results or hyperlinks

Wait for New Window

Checks for a new window to open. Used after clicking a link or button that opens a new child window in the current browser session.

Properties:

- **Max Wait Time** (optional): Maximum time to wait in seconds for the new window to appear (default: 30)
- **Window Identifier** (output): Variable name storing the identifier for the new window (default: `window_id`)

Usage Notes:

- If time exceeds and window not found, an exception occurs
- Window identifier can be used to switch focus or close the window

Validation:

- Max wait time must be an integer or decimal

Switch Window

Changes browser windows. After switching focus to the specified window, subsequent browser actions interact with that window.

Properties:

- **Window Identifier** (optional): The window you want to switch to, typically stored in a variable (e.g., `window_id`). Leave empty to switch to the main window

Close Window

Closes a browser window. Once closed, the automation returns to the main window of the current browser session.

Properties:

- **Window Identifier** (required): The window you want to close, typically stored in a variable (e.g., `window_id`)

Save to Clipboard

Copies text to the clipboard. Used to save text that you can paste later on. Existing clipboard text is replaced.

Properties:

- **Text to Save** (required): The text you want to copy to the clipboard (e.g., "Order #12345")

Paste from Clipboard

Inserts text from the clipboard. Used to paste copied text into webpage fields.

Properties:

- **Find Element with Description** (radio button): Use natural language description to identify the element with AI (default: ON)
 - **Description Field** (conditionally required): Natural language description (e.g., "Order number field")

- **Find Element with Selector** (radio button): Use XPath expression to identify the element (default: OFF)
 - **XPath Field** (conditionally required): XPath expression (e.g., "//input[@id='xyz']")
- **Self-healing** (checkbox): Use AI to handle unexpected popups automatically (default: OFF, optional)
- **Replace Existing Text** (checkbox): Clear any pre-existing text in the field before pasting (default: ON)

Element Identification:

- At least one element identification method must be enabled

Refresh Webpage

Reloads the current webpage. Some pages may redirect on refresh.

Properties:

- **Action Title** (text-only): Display name for the action

Restart Browser

Closes and reopens the browser. A new session starts from a blank page.

Properties:

- **Restart All** (checkbox): Close and restart all browser windows, not just the current one (default: OFF, optional)

Limitations

Element Identification Uncertainty: If the AI cannot reliably identify UI elements, actions fail with 'LowConfidence' errors. In these cases, use more specific XPath selectors or element descriptions.

Browser Compatibility: The system operates exclusively with Chrome browser instances and utilizes Playwright as the underlying automation framework, which means compatibility is limited to Chrome-supported web technologies and may not work with browser-specific features from other vendors. **Dynamic Content Challenges:** Highly dynamic pages with frequent layout changes,

heavy JavaScript frameworks, or complex single-page applications may require additional wait times, retry logic, or specialized handling approaches. **Performance Overhead:** AI-powered features like natural language element targeting and self-healing capabilities require more processing time than direct XPath targeting.

Excel

Excel actions enable you to read, update, and create Excel files in your automations. All Excel actions support both .xlsx and .xlsm file formats. Workbook identifiers maintain references across multiple operations.

Note

While values can be read from .xlsm files, macro execution is not supported.

Create New Workbook

Creates a blank Excel file. Outputs a workbook identifier you can use to update the spreadsheet in subsequent steps.

Properties:

- **Workbook Identifier** (output): Variable name storing the workbook reference for future actions (default: `workbook_id`)

Open Existing Workbook

Loads an existing Excel file. Returns a workbook identifier you can use to update the spreadsheet in subsequent steps. Supports only Excel (.xlsx) files.

Properties:

- **Excel File** (required): The .xlsx file to open, typically stored in a variable (e.g., `my_file`)
- **Workbook Identifier** (output): Variable name storing the workbook reference (default: `workbook_id`)

File Requirements:

- Must be a valid Excel (.xlsx) file
- File must be accessible as a media file object

Save Workbook

Saves updates to an Excel file. Allows you to update the file name and file type for the saved workbook.

Properties:

- **File Name** (required): Name for the saved file without extension (e.g., "Monthly Report")
- **File Type** (dropdown): Output format - currently supports XLSX (default: XLSX)
- **Saved File** (output): Variable storing the saved file object (default: saved_file)

Read Sheet

Gets data from a range of cells. The action stores the output in a data table variable.

Properties:

- **Workbook Identifier** (required): The workbook to read from (e.g., workbook_id)
- **Sheet Name** (required): Worksheet or tab name (default: "Sheet1")
- **Cell Range** (optional): Range specification - supports multiple formats:
 - Starting cell: "A2" (reads all data from anchor point)
 - Exact range: "A1:B10"
 - Column range: "A:B"
 - Row range: "1:3"
 - Empty: reads entire sheet
- **Include Headers** (checkbox): Treats first row as column headers when enabled. When disabled, uses default naming (Column0, Column1, etc.) (default: TRUE)
- **Data Table** (output): Variable storing the extracted data (default: excel_table)

Formula Handling:

- The action automatically calculates formulas and stores results in the data table.

Read Cell

Gets the value from a cell. Used to read an individual cell value from a worksheet.

Properties:

- **Workbook Identifier** (required): The workbook to read from (e.g., `workbook_id`)
- **Sheet Name** (required): Worksheet name (default: "Sheet1")
- **Cell Reference** (required): Cell location (e.g., "A1")
- **Cell Value** (output): Variable storing the cell content (default: `cell_value`)

Write to Sheet

Outputs a data table to a sheet. Used to write a range of rows and columns to a worksheet.

Properties:

- **Data Table** (required): The data table to write (e.g., `my_table`)
- **Workbook Identifier** (required): Target workbook (e.g., `workbook_id`)
- **Sheet Name** (required): Target worksheet (default: "Sheet1")
- **Start at Cell** (optional): Starting position for data placement (default: "A1")
- **Include Headers** (checkbox): Writes column headers when enabled (default: TRUE)

Write to Cell

Outputs a value to a cell. Used to update individual cells in a worksheet.

Properties:

- **Value to Write** (required): Content for the cell (e.g., "Order #12345")
- **Workbook Identifier** (required): Target workbook (e.g., `workbook_id`)
- **Sheet Name** (required): Target worksheet (default: "Sheet1")
- **Cell Reference** (required): Target cell location (e.g., "A1")

Write New Row

Adds a row of data to the sheet. The new row is appended to the end of the existing data.

Properties:

- **Row Values** (required): Array of values for the new row, starting from first column (e.g., ["Q1", "Sales", 100])
- **Workbook Identifier** (required): Target workbook (e.g., `workbook_id`)
- **Sheet Name** (required): Target worksheet (default: "Sheet1")

Data Validation:

- Must provide values as an array format
- Values are written sequentially starting from the first column

Create New Sheet

Adds a blank worksheet. The new sheet is added to the end of the workbook.

Properties:

- **Workbook Identifier** (required): Target workbook (e.g., `workbook_id`)
- **Sheet Name** (required): Name for the new worksheet (e.g., "Sales Data")

Copy Sheet

Creates a copy of the worksheet. The new sheet is created within the same workbook.

Properties:

- **Workbook Identifier** (required): Target workbook (e.g., `workbook_id`)
- **Sheet Name to Copy** (required): Source worksheet name (e.g., "Original Sheet")
- **New Sheet Name** (required): Name for the duplicated sheet (e.g., "Original Sheet (Copy)")

Rename Sheet

Updates the name of an existing worksheet.

Properties:

- **Workbook Identifier** (required): Target workbook (e.g., `workbook_id`)

- **Current Sheet Name** (required): Existing worksheet name (e.g., "Sheet1")
- **Updated Sheet Name** (required): New name for the worksheet (e.g., "Q1 Data")

Delete Sheet

Removes a sheet from a workbook. Cannot delete the last remaining sheet of a workbook.

Properties:

- **Workbook Identifier** (required): Target workbook (e.g., `workbook_id`)
- **Sheet Name to Delete** (required): Worksheet to remove (e.g., "Sheet1")

Constraints:

- Cannot delete the last remaining sheet in a workbook

Set Cell Color

Updates cell background color. Used to highlight a specific range of cells.

Properties:

- **Cell Color** (required): RGB hex code format (e.g., "FF0000" for red)
- **Workbook Identifier** (required): Target workbook (e.g., `workbook_id`)
- **Sheet Name** (required): Target worksheet (e.g., "Sheet1")
- **Cell Range** (required): Range to format - supports:
 - Single cell: "A1"
 - Exact range: "A1:B10"
 - Column range: "A:B"
 - Row range: "1:3"

Get Cell Color

Reads the cell background color. Outputs the color in RGB hex code format (e.g., "FF0000" for red).

Properties:

- **Workbook Identifier** (required): Source workbook (e.g., `workbook_id`)
- **Sheet Name** (required): Source worksheet (e.g., "Sheet1")
- **Cell Reference** (required): Cell to read (e.g., "A1")
- **Cell Color** (output): Variable storing RGB hex code (default: `cell_color`)

Hide Rows

Makes rows hidden in the sheet. The data remains intact but is not visible.

Properties:

- **Workbook Identifier** (required): Target workbook (e.g., `workbook_id`)
- **Sheet Name** (required): Target worksheet (e.g., "Sheet1")
- **Start Row** (required): First row number to hide (Excel numbering starts at 1)
- **End Row** (optional): Last row number to hide. If empty, only the start row is hidden

Row Numbering:

- Excel row numbers start at 1 (not 0)
- Must provide integer values

Unhide Rows

Makes hidden rows visible. Used to show previously hidden rows.

Properties:

- **Workbook Identifier** (required): Target workbook (e.g., `workbook_id`)
- **Sheet Name** (required): Target worksheet (e.g., "Sheet1")
- **Start Row** (required): First row number to unhide (Excel numbering starts at 1)
- **End Row** (optional): Last row number to unhide. If empty, only the start row is made visible

Best practices and limitations

Workbook Identifier Management

- Store workbook identifiers in descriptive variables (e.g., `sales_workbook`, `report_file`)

- Reuse the same identifier across multiple actions on the same workbook
- Always create or open a workbook before performing data operations

Range Specifications

- Use exact ranges ("A1:B10") for precise data operations
- Use column ranges ("A:B") when working with entire columns
- Use row ranges ("1:3") for header or summary operations
- Leave range empty to process entire sheets

Performance Optimization

- Read entire ranges when possible instead of individual cells
- Batch write operations using data tables rather than individual cell writes
- Save workbooks only when all modifications are complete

Limitation

File compatibility is limited to modern Excel (.xlsx) format only - legacy Excel (.xls) files are not supported.

Data tables

Data table actions enable you to work with structured data in table format. These actions allow you to create, transform, and output tabular data in your automations.

Sort Table

Orders a table by a column. Used to organize your data in ascending or descending order.

Properties:

- **Data Table** (required): The table variable to sort (e.g., my_table)
- **Column Name to Sort** (required): Name of the column to sort by (e.g., "Total amount")
- **Sort Order** (dropdown): Choose "Ascending" (smallest first) or "Descending" (largest first) - default: Descending
- **Sorted Table** (output): Variable name for the new sorted table

Filter Table

Keeps rows matching a criteria. Used to extract relevant rows from a larger data set.

Properties:

- **Data Table** (required): The table variable to filter (e.g., `my_table`)
- **Filter Expression** (required): Boolean expression using column names and operators (`==`, `>`, `<`, `!=`). Combine conditions with `&` (AND) or `|` (OR). Use single quotes for column names with spaces and text values (e.g., `"Team" == 'Sales' & 'Total amount' > 100`)
- **Filtered Table** (output): Variable name for the filtered table

Example filters:

```
# Column 'amount' is greater than 25
"amount > 25"

# Column 'team' equals text 'Sales'
"team == 'Sales'"

# Multiple conditions
"amount > 25 & status == 'active'"

# Grouped conditions
"(amount > 25 & team == 'Sales') | (amount > 50 & team == 'Marketing')"
```

```
# Column 'title' contains text 'Director'
"title.str.contains('Director')"
```

```
# Column 'Customer No' equals a variable customer_id
"`Customer No` == " + str(customer_id)
```

```
# Column 'ProfitCenter' equals variable profitCenter
"`ProfitCenter` == '" + str(profitCenter) + '"
```

Lookup Value

Searches a value in a table. Used to lookup a value in one column and get the corresponding value from another column in the same row.

Properties:

- **Data Table** (required): The table to search in (e.g., `my_table`)
- **Column Name to Search** (required): Column containing the lookup value (e.g., "Employee ID")
- **Value to Search For** (required): The value to find (e.g., "12345")
- **Column Name to Output** (required): Column to retrieve the result from (e.g., "Date of hire")
- **Cell Value** (output): Variable storing the found value. Returns first match or empty if not found.

Add Columns

Creates new columns in a table. Newly added columns are appended to the end of the existing table.

Properties:

- **Data Table** (required): The table to modify (e.g., `my_table`)
- **Column Names to Add** (required): Array of new column names (e.g., ["Name", "Address"])
- **Default Value** (optional): Initial value for all cells in new columns (e.g., "N/A")
- **Updated Table** (output): Variable name for the modified table

Remove Columns

Deletes columns from a table. Outputs a table with all of the remaining columns.

Properties:

- **Data Table** (required): The table to modify (e.g., `my_table`)
- **Columns to Remove** (required): Array of column names or index numbers. Index numbers start at 0 and can be specific numbers (e.g., [0,1,2]) or ranges (e.g., `range(0,2)`)
- **Updated Table** (output): Variable name for the modified table

Keep Columns

Drops extra columns from a table. Used to select a specific subset of columns you want to keep.

Properties:

- **Data Table** (required): The table to modify (e.g., `my_table`)
- **Columns to Keep** (required): Array of column names to retain (e.g., ["Name", "Address"])

- **Updated Table** (output): Variable name for the modified table

Add New Row

Appends a new row to a table. The new row can be created with specific values or as a blank row and is added to the bottom of the table.

Properties:

- **Data Table** (required): The table to modify (e.g., `my_table`)
- **Row Values** (optional): Array of values for the new row, starting from the first column (e.g., `["Q1", "Sales", 100]`). If empty, adds a blank row. Missing values result in blank cells.
- **Updated Table** (output): Variable name for the modified table

Remove Rows

Deletes rows from a table. Outputs a table with all of the remaining rows.

Properties:

- **Data Table** (required): The table to modify (e.g., `my_table`)
- **Rows to Remove** (required): Array of row positions (0-based indexing). Index numbers start at 0 and can be specific numbers (e.g., `[0,1,2]`) or ranges (e.g., `range(0,2)`)
- **Updated Table** (output): Variable name for the modified table

Remove Duplicates

Deletes duplicate rows. Used to create a dataset of unique rows based on specific columns.

Properties:

- **Data Table** (required): The table to clean (e.g., `my_table`)
- **Columns to Check** (optional): Array of column names for duplicate detection (e.g., `["Name", "Address"]`). If empty, checks entire rows for uniqueness. Duplicates are identified by combined values across specified columns.
- **Duplicate Row to Keep** (dropdown): Choose "First" or "Last" occurrence to retain (default: First)
- **Updated Table** (output): Variable name for the modified table

Append Tables

Combines the rows of two tables. Used to add data from one table to another.

Properties:

- **Table to Append To** (required): Main table receiving additional rows (e.g., `main_table`)
- **Table to Add** (required): Source table providing rows to append (e.g., `new_data`)
- **Handle Column Differences** (dropdown):
 - "Add": Keep all columns from both tables
 - "Ignore": Only keep columns matching the primary table
 - "Error": Require exact column matches
- **Combined Table** (output): Variable name for the merged table

Create New Table

Creates an empty table. Used to set up a table with required columns to add rows to in subsequent steps. The new table has no rows.

Properties:

- **Column Names** (optional): Array of column names for the new table (e.g., `["Name", "Address"]`). If empty, creates a table with no columns.
- **New Table** (output): Variable name for the newly created table

Convert Text to Table

Transforms delimited text into a structured table.

Properties:

- **Text to Convert** (required): Delimited text containing table data (e.g., `"Year,Qty 2001,100"`)
- **Value Separator** (optional): Character separating values in rows (default: `","`)
- **Newline Separator** (optional): Character separating rows (default: `" "`)
- **Has Headers** (checkbox): Whether first row contains column names. If True, uses first row as headers; if False, generates default names (`Column0`, `Column1`, etc.)
- **New Table** (output): Variable name for the newly created table

Convert Table to HTML

Creates an HTML formatted table. Used to output your table as formatted text for a document, email, and more.

Properties:

- **Data Table** (required): The table to format (e.g., `my_table`)
- **Formatted Table Text** (output): Variable storing the HTML formatted text

Examples

Loop through rows in a table

To loop through rows in a data table, use the "Loop through items" action under "Process flow". Provide the data table variable as an input for "Collection of items". Update the "Item reference" to `row` as each item represents a single row of the table.

Use or update cell values in a row

When looping through rows in a table, refer to individual cell values using the syntax of `row["column name"]` where column name is replaced with the name of the column for the specific cell.

To update the value of a cell in a row, use the "Save value" action under "General". The "Value to save" is the value you want to update the cell with. The "Variable name" is the reference to the cell using the same syntax as above, `row["column name"]`.

Orchestration

Orchestration enables you to automate complex end-to-end processes, coordinating automated tasks with human decisions where needed. The orchestration capabilities in Amazon Quick Automate include case management for tracking work items and human-in-the-loop integration for incorporating human judgment.

Case management

Cases track individual items of work for your automation to process such as a new claim, an IT support ticket, a customer complaint, or a new invoice. Each case has:

- A unique reference name (e.g., "Order-12345")
- A case type for grouping related work (e.g., "Sales orders")
- Custom data stored as key-value pairs
- A defined lifecycle with status tracking and an audit trail
- Associated human-in-the-loop tasks if applicable

Case status lifecycle

To track the progress of work through your automation, cases follow a defined lifecycle with specific status transitions:

- **Ready** - The case has been created and is waiting to be picked up for processing. This is the initial state of all new cases.
- **In Progress** - The case is actively being processed by the automation. You can monitor real-time progress through the logs.
- **Pending Resolution** - The case is waiting for a human-in-the-loop task to be completed before processing can continue. Cases automatically move back to Ready status after tasks are resolved.
- **Successful** - The case has completed all processing steps without any exceptions.
- **Failed** - The case encountered one of the following exception types:
 - **Business Exception** - The case encountered a handled business rule violation and stopped processing. Detailed exception information is available in the logs.
 - **System Exception** - The case experienced a technical error and stopped processing. Detailed error information is available in the logs.

Actions

Amazon Quick Automate provides the following actions to create, process, and manage cases throughout their lifecycle.

Create new case

Creates a single case to track work. Each case has a reference name (e.g., "Order-12345") and stores custom data as key-value pairs. Cases are grouped by type (e.g., "Sales orders") for organization and reporting. Properties:

- **Title (optional)** - Action name shown in the process visualization (e.g., "Create sales order case")
- **Reference name (required)** - Business identifier used for searching and tracking cases in monitoring dashboards (e.g., "Order-12345")
- **Case type (required)** - Label that groups related cases by business purpose. Provides organization and metric tracking (e.g., "Sales orders", "Support tickets")
- **Custom data (required)** - Additional case information stored as key-value pairs. This data is visible to users and can be retrieved, updated, or appended to during case processing (e.g., {"Product": "ABC-12345", "Qty": "100", "Priority": "High"})

Create multiple cases

Bulk creates cases from a data table. Each row becomes a new case, inheriting the specified case type and using a designated column for reference names. All other columns become custom data fields. Properties:

- **Title (optional)** - Action name shown in the process visualization (e.g., "Create cases from invoice batch")
- **Case type (required)** - Label that groups related cases by business purpose (e.g., "Invoice processing")
- **Cases data table (required)** - Data table where each row creates a new case. All columns are stored as custom data (e.g., cases_table)
- **Reference column name (required)** - Name of the column containing the reference name identifier for each case (e.g., "Invoice number")

Process cases

Loops through cases in 'Ready' status. Automatically manages case status by moving cases to 'In Progress' during processing and marking them as successful or failed based on completion. When adding actions inside the Process cases loop, uncaught exceptions are handled for you, marking the case as failed with the associated error details and moving to the next case. The action also provides access to case data during processing. Properties:

- **Title (optional)** - Action name shown in the process visualization (e.g., "Process pending invoices")
- **Case type (required)** - Type of cases to process (e.g., "Invoice processing")
- **Current case (output)** - Variable containing the case being processed in each loop. Access case data using `current_case["custom_data"]["key_name"]`

Update case data

Saves changes to case information during processing. Can add new fields or modify existing data. Only available while a case is 'In Progress'. Used to track processing results and manage case lifecycle. Properties:

- `Title` (optional) - Action name shown in the process visualization (e.g., "Update case status")
- `Case ID` (required) - Unique identifier of the case to update (e.g., `current_case["case_id"]`)
- `Data to update` (required) - New or modified information to store as key-value pairs (e.g., `{"Status": "Approved", "ProcessedDate": "2024-01-20"}`)
- `Updated case` (output) - Variable containing the modified case. Access updated data using `updated_case["custom_data"]["key_name"]`

Search cases

Retrieves cases for reporting and analysis. Filter by case type, status, or reference name. Returns results as a data table containing case information and custom data. Properties:

- `Title` (optional) - Action name shown in the process visualization (e.g., "Find overdue cases")
- `Case type` (required) - Type of cases to search (e.g., "Support tickets")
- `Filter by status` (optional) - Filter for cases based on their current status (e.g., "Pending Resolution")
- `Filter by reference name` (optional) - Filter for cases with a specific reference name (e.g., "Ticket-12345")
- `Maximum results` (optional) - Limits the number of cases to return. Leave empty to return all cases up to 10,000 (e.g., 100)
- `Search results` (output) - Variable name that will store the data table of cases retrieved by the search (e.g., `cases_table`)

Human-in-the-Loop (HITL)

While many decisions can be automated, some scenarios require human judgment. Amazon Quick Automate's HITL capabilities allow you to seamlessly incorporate human decision-making into your automated processes. There are two ways to create HITL tasks:

- Using the Create user task action
- Using HITL chat through a Custom agent

Create user task

Creates a task for human input. Tasks appear in a list for users to review, make decisions, or provide information. Tasks can only be created when in a case processing loop while the case is in progress. Creating a task moves the case to 'Pending task' status and the automation continues with the next case. Once the task is resolved, the case returns to 'Ready' status for further processing. Properties:

- **Title (required)** - Subject line that appears in the user's task list. Should clearly describe what needs to be done (e.g., "Review loan application", "Approve expense report")
- **Description (required)** - Detailed instructions and context provided to users completing the task (e.g., "Please review the attached application details and either approve or reject based on the following criteria...")
- **Form (required)** - Interface for collecting structured input from users
- **Severity (required)** - Priority level to help users triage tasks:
 - **Critical** - Urgent attention required, highest priority
 - **Moderate** - Normal priority (default)
 - **Low** - Non-urgent task
- **Due in (required):**
 - **Number** - Quantity of time units (e.g., 24)
 - **Unit** - Hours or Days
- **Case ID (required)** - Unique identifier of the case this task is for (e.g., `current_case["case_id"]`)

Important

When creating a HITL task:

- The Process cases loop automatically updates the case status to Pending Resolution and moves on to the next case
- Do not add actions after creating a task inside the Process cases loop as they won't execute for the current case

Building forms

When creating a HITL task, you can build custom forms to collect exactly the information needed from human reviewers. The following components are available for form design:

- Text field - Single or multi-line text input for collecting user inputs (e.g., "Please provide reason for rejection in detail")
- Dropdown - Selection from predefined options (e.g., ["Approve", "Reject", "Escalate"])
- Date/time picker - Calendar-based input
- Media - Display images or documents for review
- Display text - Read-only information to provide context (e.g., case details, policy references)

HITL chat

For scenarios requiring more dynamic interaction, Amazon Quick Automate also supports HITL chat through Custom agents. This enables natural language conversation between the agent and human reviewer. The agent can create a HITL task when it needs human input. The case will remain In progress while the agent chats with the human reviewer. Once sufficient input is provided, the agent will continue.

```
# Example agent prompt for HITL chat
"""
Review the data and if any values are outside normal ranges,
create a HITL task for human review. Include the specific
values and acceptable ranges in the task description.
"""
```

Note

HITL chat tasks will be given a default severity (Moderate) and due in (24 hours).

Task resolution and case re-entry

After a human completes a HITL task, your automation needs to handle the case's return to processing. Understanding how to work with task resolution data is key to implementing effective human-in-the-loop workflows. When a task is resolved:

- The associated case automatically returns to Ready status
- The case becomes available for processing again through Process cases
- The `latest_task_resolution` field contains all form inputs and decisions

Example checking for HITL resolution and handling the case accordingly:

```
# Example handling HITL resolution
if current_case["latest_task_resolution"]:
    # Handle case returning from HITL
    resolution_data = current_case["latest_task_resolution"]["form_data"]
    if resolution_data["decision"] == "Approve":
        # Process approved case
    else:
        # Handle rejected case
else:
    # First time processing this case
```

Best practices for handling HITL resolution:

- Always check latest_task_resolution when processing cases
- Structure your logic to handle both new cases and those returning from HITL
- Consider implementing different processing paths based on human decisions
- Include error handling for missing or invalid resolution data

Exception handling

Exception handling allows you to manage errors and unexpected conditions in your automations. Amazon Quick Automate provides capabilities for creating, raising, and handling exceptions appropriately to ensure robust error management and process continuity.

Types of exceptions

Amazon Quick Automate supports two types of exceptions:

- Business exceptions - Used for expected errors from business rule violations such as:
 - Missing required data
 - Invalid input values
 - Policy violations
- System exceptions - Used for technical or infrastructure errors such as:
 - Connection failures
 - Service timeouts
 - Authentication errors

Actions

Raise exception

Interrupts the process flow with an exception. Used to trigger or throw an exception to either end the process or to take specific actions to handle the exception scenario. If you do not want to end the process immediately, add an Exception flow to your Process step to handle the exception. Properties:

- **Title** (optional) - Action name shown in the process visualization (e.g., "Raise validation error")
- **Create custom exception** (required) - When selected on, create a new exception with custom details. When off, raise or re-raise an existing exception stored in a variable
- **Exception type** (required) - Category of the exception (e.g., "Business exception", "System exception")
- **Exception reason** (required) - Short identifier used to categorize similar exceptions for reporting (e.g., "REQUIRED_FIELD_MISSING", "INVALID_FORMAT")
- **Exception message** (required) - Detailed description that will help with troubleshooting (e.g., "Customer email is required but not provided")
- **Exception to raise** (required when Create custom exception is off) - The exception you want to raise or re-raise, typically stored as a variable (e.g., new_exception)

Create custom exception

Creates an exception without raising it. Used to prepare an exception variable that can be raised later in subsequent steps. Properties:

- **Title** (optional) - Action name shown in the process visualization (e.g., "Create validation exception")
- **Exception type** (required) - Category of the exception (e.g., "Business exception", "System exception")
- **Exception reason** (required) - Short identifier used to categorize similar exceptions for reporting (e.g., "REQUIRED_FIELD_MISSING")
- **Exception message** (required) - Detailed description that will help with troubleshooting (e.g., "Customer email is required but not provided")
- **Error code** (optional) - Numerical code used to distinguish different types of errors (e.g., 404 for not found, 503 for service unavailable)
- **Caused by** (optional) - Reference to another exception that triggered this one, typically stored as a variable (e.g., original_exception)

- **New exception** (output) - Variable that will store the newly created exception (e.g., `new_exception`)

Exception flow

Defines how to handle exceptions that occur within a Process step. Used to create an alternative path to follow when exceptions occur in your process. After an exception is handled, the process continues with the next Step. Raise the exception again if you want to end the process. Add an Exception flow by clicking the **Exception flow** plus button found in the Process step. Properties:

- **Title** (optional) - Action name shown in the process visualization (e.g., "Handle validation errors")
- **Exception reference** (required) - Variable name to refer to the exception within the Exception flow. Access exception details using `variable["property"]` where property can be: type, reason, message, code, or `caused_by` (e.g., `error`)

Note

Exception flows can only be added to Process steps. They are not found otherwise in the Actions panel.

Working with exceptions

Exception properties

When handling exceptions, you can access these properties:

- `type` - Category of the exception ("Business exception" or "System exception")
- `reason` - Short identifier for the exception
- `message` - Detailed description of the error
- `code` - Numerical error code if provided
- `caused_by` - Original exception that triggered this one if applicable

Best practices

- Use descriptive exception reasons for easy categorization and reporting
- Include detailed messages for troubleshooting
- Implement retry logic for temporary system exceptions

- Consider human intervention for business exceptions
- Consider the impact on case status when handling exceptions
- Ensure critical cleanup operations occur even after exceptions such as restarting a browser
- Monitor exception frequencies and patterns
- Review exception logs regularly
- Re-raise exceptions once handled if you do not want to proceed with the remaining steps

Files

- **Duplicate file** - Creates a copy of a file. Used to create a new version of the same file.
- **Zip files** - Creates a compressed zip file. Used to combine multiple related files into a single package.
- **Unzip file** - Extracts files from a zip. Used to unpack individual files from a single zip file.

PDF

- **Create PDF file** - Generates a PDF file from text. Used to create a new PDF document with formatted contents.
- **Read PDF text** - Extracts text from a PDF file. Used to get the content of text-based PDF documents. Does not extract text from scanned documents or preserve the structure of the extracted text.
- **Split PDF file** - Extracts pages from a PDF. Used to create a new PDF from selected pages of an original file.
- **Merge PDF files** - Combines multiple PDF files. Used to create a single PDF from several documents by appending their pages.
- **Get page count** - Gets the count of pages in a PDF. The number of pages can be used in subsequent steps to extract text from specific pages or split the PDF.

Code

Code actions in Quick Automate let you implement custom logic using Python code blocks, going beyond what standard automation actions support. They are well-suited for complex data transformations and calculations, and run within a restricted Python environment to maintain security.

Two Types of Code Actions:

- **Single Line Expressions:** Quick, one-line operations that modify variables without returning values
- **Custom Code Blocks:** Multi-line Python functions for complex logic with parameters and return values

When to use code actions:

Use code blocks when you need to:

- Perform complex data transformations not available in standard actions
- Implement custom business logic or calculations
- Process or manipulate data structures (lists, dictionaries, JSON)
- Work with dates, times, and time zones in custom ways
- Parse or format strings with complex patterns
- Optimize automation performance by consolidating multiple actions into one code block

When not to use code actions:

Avoid code blocks when:

- A standard automation action already exists for your use case
- The operation is simple enough for built-in actions
- You need to interact with external APIs (use REST API integration instead)
- You need to access file systems or databases (use appropriate integrations)

How to Access code actions:

Code blocks are available through multiple interfaces:

- **Actions Panel (Recommended):**
 - Open your automation in the automation builder
 - Click on the Actions Panel on the right side
 - Locate "Custom Code Block" under the Code Actions section
 - Drag and drop the code block into your automation workflow

- **Build with Assistant:**

- Available during plan generation when the assistant determines custom code is needed
- The assistant automatically suggests code blocks for complex operations
- You can request code blocks by describing your custom logic needs

Actions available:

Single Line Expressions

Single Line Expressions execute one-line Python statements that perform operations without returning a value. They are ideal for quick modifications to existing variables like appending to lists, updating dictionaries, or performing simple calculations that modify state.

Properties:

- **Expression (required):** The Python expression to execute (e.g., "my_list.append('new item')")

Examples:

- **Appending to list**

```
my_list.append("1")
my_list.append(new_item)
```

- **Removing List Items**

```
task_list.remove(completed_task)
```

Custom Code Block

Custom Code Blocks are multi-line Python functions that execute complex logic, accept parameters, and return values. They are the fallback option when standard automation actions and Single Line Expressions are not sufficient for your needs.

Properties:

- **Function Title (required):** Name identifier for the code block (e.g., "Calculate_Total")
- **Function (required):** Python code block that contains your custom logic.

- **Step 1: Define Parameters**
 - Click the "Edit" button to open the code editor
 - In the Parameter panel, click "Add" to create new parameters
 - Enter parameter names that match your automation variables
 - Parameters are available as function arguments
- **Step 2: Write Your Python Code**
- **Return Value (optional):** Variable name to store the function's output
 - Follow the required code block structure (see below)
 - Implement your custom logic within the function
 - Use only approved libraries and built-in functions
 - Include a return statement if you need to output data

Code Block Structure

All code blocks must follow this specific format:

```
@code_block()
def your_function_name(parameter1, parameter2, parameter3):
    -----
    """
    Optional: Add a docstring describing what your function does
    """
    # Your custom logic here
    result = parameter1 + parameter2 + parameter3

    return result
```

Built-in Python Functions and Imports

All standard Python built-in functions are available without imports (len, str, int, etc.)

Approved Standard Libraries (Security Restricted)

Code blocks can ONLY import these standard libraries:

- `base64` - Base64 encoding/decoding
- `datetime` - Date and time operations

- `json` - JSON parsing and generation
- `math` - Mathematical functions
- `re` - Regular expressions
- `zoneinfo` - Time zone handling

Note

- Cannot import standard libraries other than the ones listed above.
- Cannot install or import third-party libraries. `pip install` is not supported in code blocks.

Limitations

- Restricted Python environment with limited library access. The execution environment is based on RestrictedPython, a subset of Python 3.10.
- Code blocks cannot invoke other code blocks or actions

Best Practices

- Keep code blocks simple and focused
- Use descriptive function names
- Always prefer pre-built actions when available
- Test your code blocks thoroughly as focused debugging options are limited)

Example Use Cases

- Mathematical operations (Calculate circle properties - radius as parameter)

```
def function
(radius):

    return {
        "radius": radius,
        "diameter": 2 * radius,
        "circumference": round(2 * math.pi * radius, 2),
```

```
"area": round(math.pi * radius ** 2, 2)
}
```

- Getting current date/time

```
def function
():

    now = datetime.datetime.now()

    return {
        "current_date": now.strftime("%Y-%m-%d"),
        "current_time": now.strftime("%H:%M:%S"),
        "formatted": now.strftime("%B %d, %Y at %I:%M %p"),
        "iso_format": now.isoformat(),
        "timestamp": now.timestamp()
    }
```

- Calculating date differences - Start date and end date as parameters

```
def function
_(start_date_str, end_date_str):

    # Parse date strings (format: YYYY-MM-DD)_
    start = datetime.datetime.strptime(start_date_str, "%Y-%m-%d")
    end = datetime.datetime.strptime(end_date_str, "%Y-%m-%d")

    _# Calculate difference_
    difference = end - start
    return {
        "days": difference.days,
        "weeks": difference.days // 7,
        "start": start_date_str,
        "end": end_date_str    }
```

- Pattern matching and text manipulation using regular expressions (Validating email addresses, phone numbers, etc.)

```
def function
(email, phone, zip_code):

    email_pattern = r'^[a-zA-Z0-9._%+-]+@[a-zA-Z0-9.-]+\.[a-zA-Z]{2,}$'
    phone_pattern = r'^\d{3}-\d{3}-\d{4}$'
```



```
zip_pattern = r'^\d{5}(-\d{4})?$'

return {
    "email_valid": bool(re.match(email_pattern, email)),
    "phone_valid": bool(re.match(phone_pattern, phone)),
    "zip_valid": bool(re.match(zip_pattern, zip_code))
}
```

- **List Operations (Filter and Transform List)**

```
def function
(numbers, threshold):

    # Filter numbers above threshold and calculate statistics
    filtered = [n for n in numbers if n > threshold]

    if filtered:
        return {
            "filtered_numbers": filtered,
            "count": len(filtered),
            "sum": sum(filtered),
            "average": sum(filtered) / len(filtered),
            "min": min(filtered),
            "max": max(filtered)
        }
    else:
        return {
            "filtered_numbers": [],
            "count": 0,
            "message": "No numbers above threshold"
        }
```

Integrations

AWS Services

Amazon S3

Amazon S3 actions allow you to store and retrieve files and data.

Available Actions

- **Download File** - Downloads a file from Amazon S3 and returns a file object handle.

- **Upload File** - Uploads a file object handle to Amazon S3.
- **Copy Object** - Creates a copy of an object that is already stored in Amazon S3 from one source file to a destination location.
- **Delete Object** - Deletes a specified object from an S3 bucket.
- **List Objects** - Lists all of the objects in the given bucket. Can be filtered using a prefix.

Amazon Bedrock

Amazon Bedrock actions allow you to connect to foundation models (FMs) from leading AI companies through a single API.

- **Basic model converse** - Sends a prompt to an AI model. Use natural language to generate text, analyze content, and get AI-powered responses.
- **Advanced model converse** - Prompts model with more options. Advanced options include custom message formats, tool integrations, and granular control over model behavior.

Amazon Bedrock Agents

Amazon Bedrock Agent actions allows you to invoke AI assistants that can connect to your data sources, APIs, and backend systems.

- **Basic agent invoke** - Sends a request to an AI agent. Use for simple interactions that don't require file attachments, user input, or custom configurations.
- **Advanced agent invoke** - Calls an agent with more options. Used for advanced interactions that require file attachments, user input, or custom configurations.

Amazon Bedrock Data Automation

Bedrock Data Automation (BDA) simplifies the process of extracting valuable insights from unstructured content using generative AI.

Available Actions

- **Invoke Data Automation Async** - Process data with a data automation project or blueprints. Starts asynchronous data processing and returns an invocation ID.
- **Get Data Automation Status** - Gets details about a data automation invocation status and output location.

Amazon Textract

Amazon Textract helps you add document text detection and analysis to your automations.

- **Basic document query** - Extracts data based on questions. Each query contains the question you want to ask and the alias you want to associate. Textract provides the text answer to each question and a confidence score.
- **Basic document analysis** - Extracts data from documents. Automatically detects and extracts forms (key-value pairs), tables (structured data), and text while preserving document layout and data relationships.

Amazon Comprehend

Amazon Comprehend is a natural language processing (NLP) service for gaining insight into document content. It helps analyze text to extract key phrases, identify sentiment, and classify documents.

Available Actions

- **Classify Document** - Creates a classification request to analyze a single document in real-time using a custom model endpoint. Supports text input or document files (PDF, Word, or image).
- **Detect Key Phrases** - Detects the key noun phrases found in the text, identifying important concepts and topics within the content.

External Applications

Microsoft Outlook

The Microsoft Outlook actions allow your automations to manage emails and calendars.

- **Email**
 - **Send User Email** - Send an email message on behalf of a specific user.
 - **Reply to Email** - Reply to the sender of a message.
 - **List Folder Messages** - Get the messages in a specific folder of the signed-in user's mailbox.
 - **List User Mails** - Get the messages in the signed-in user's mailbox (including Deleted Items and Clutter folders).
 - **View Email** - Retrieve the properties and relationships of a specific message object.
 - **Update Email** - Update the properties of a message object.

- **Delete Email** - Delete a message in the specified user's mailbox.
- **List Email Attachments** - Retrieve a list of attachment objects attached to a message.
- **Get Attachment** - Read the properties, relationships, or raw contents of an attachment.
- **Folders**
 - **List Mail Folders** - Get the mail folder collection directly under the root folder of the signed-in user.
 - **Move Email to Folder** - Move a message to another folder within the specified user's mailbox.
 - **Get Mailbox Settings** - Get the user's mailbox settings, such as automatic replies, date format, and time zone.
- **Calendar**
 - **List Calendar Events** - Fetch a list of calendar events for the specified user or resource from Microsoft Outlook.
 - **Create Calendar Event** - Create an event in the user's default calendar or specified calendar.
 - **Update Calendar Event** - Update the properties of an existing event object.
- **Other**
 - **List Places** - Get a collection of place objects (rooms or room lists) defined in the tenant.
 - **List Users** - Retrieve a list of users in the organization from Microsoft Outlook with their basic information.
 - **List Contacts** - Get a contact collection from the default contacts folder of the signed-in user.

Microsoft Teams

The Microsoft Teams actions allow your automations to communicate with team members through channels and chat messages, as well as manage meetings.

- **Messages**
 - **List Chats** - View all chat conversations
 - **Create Chat** - Start a new chat conversation
 - **Get Individual Chat** - View details of a specific chat
 - **List All Chat Messages** - View messages in a chat conversation
 - **Send Chat Message** - Send a new message in a chat
- **Teams**
 - **List Teams** - View all teams accessible to the authenticated user

- **Get Team** - View detailed information about a specific team
- **List All Team Members** - View members of a team
- **Invite User To Team** - Add a new member to a team
- **Channels**
 - **List All Channels** - View all channels within a specific team
 - **Create Channel** - Create a new channel within a team
 - **Get Channel** - View detailed information about a specific channel
 - **List All Channel Messages** - View messages posted in a channel
 - **Send Channel Message** - Post a new message to a channel
 - **Reply To Channel Message** - Add a reply to an existing channel message
 - **Invite Channel Member** - Add a member to a specific channel
- **Meetings**
 - **Create Online Teams Meeting** - Schedule a new online meeting
 - **Get Online Teams Meeting** - View details of a specific meeting
 - **List All Recordings** - View recordings from a meeting
 - **List All Transcripts** - View transcripts from a meeting
- **Users**
 - **List Users** - View list of users in the organization

Microsoft SharePoint

The Microsoft SharePoint allows you to centrally store, manage, and share documents and information through your automations.

- **Lists and items**
 - **Get List** - Get a list of operations associated with a list
 - **Create Item** - Create a new list Item in a list
 - **Update Item** - Update the properties on a listItem
 - **Delete Item** - Removes an item from a list
 - **Get Item** - Returns the metadata for an item in a list
 - **List Item** - Get the collection of items in a list
- **Sheets and files**

- **Add Sheet** - Add a new worksheet to the workbook
- **Delete Sheet** - Remove worksheet from workbook
- **Update Sheet** - Edit worksheet properties (rename, visibility)
- **List Sheets** - Retrieve a list of worksheet objects
- **Read Sheet** - Retrieve properties of a worksheet
- **Upload File** - Creates or updates a workbook file
- **Ranges and cells**
 - **Read Range** - Retrieves values from a specified range
 - **Delete Range** - Deletes cells in a specified range
 - **Write Range** - Updates values in a specified range
 - **Get Used Range** - Returns the used range of the given worksheet
 - **Clear Range** - Clear range values including format, fill, and border
 - **Read Cell** - Gets the value from a specific cell
 - **Write Cell** - Sets the value of a specific cell

Microsoft OneDrive

Microsoft OneDrive actions allow you to store, sync, and share files. You can interact directly with Microsoft Excel files saved to OneDrive for spreadsheet actions.

- **File management**
 - **Get Item** - View details of a specific file or folder
 - **Get Drive** - View OneDrive account details and storage information
 - **List Items** - View all files and folders in a drive
 - **Create Item / Create Folder** - Create a new folder or item in OneDrive
 - **List Children** - View items contained within a specific folder
 - **Delete Item** - Remove a file or folder
 - **Update Item / Move Item** - Update metadata or move files between folders
 - **Copy Item** - Create a copy of a file or folder in a different location
 - **Add Permissions** - Configure sharing and access permissions for files
 - **Upload File** - Add new files to OneDrive

- **Add Sheet** - Create a new worksheet in an Excel workbook
- **Clear Range** - Remove content and formatting from a range of cells
- **Delete Range** - Remove cells, rows, or columns from a worksheet
- **Delete Sheet** - Remove a worksheet from a workbook
- **List Sheets** - View all worksheets in a workbook
- **Read Cell** - Get the value of a specific cell
- **Read Range** - Get values from a range of cells
- **Read Sheet** - View contents of an entire worksheet
- **Read Used Range** - Get the range of cells that contain data
- **Update Sheet** - Modify worksheet properties
- **Write Cell** - Set the value of a specific cell
- **Write Range** - Set values for a range of cells

Salesforce

Salesforce actions allow you to automate the process of managing customer relationships across cases, opportunities, leads, campaigns, contacts and more.

- **Cases**
 - **Get Case List** - Gets a list of customer support cases from Salesforce, including their metadata and recent items.
 - **Get Case** - View details of a specific case
 - **Create Case** - Create a new support case
 - **Delete Case** - Remove an existing case
 - **Update Case** - Edit case details and information
- **Opportunities**
 - **Get Opportunities** - View list of sales opportunities
 - **Get Specific Opportunity** - View details of a specific opportunity
 - **Create Opportunity** - Create a new sales opportunity
 - **Update Opportunity** - Edit opportunity details
 - **Delete Opportunity** - Remove an existing opportunity

- **Update Lead** - Edit lead details and information
- **Campaigns**
 - **Update Campaign** - Edit campaign details and information
- **Feed Items**
 - **Update Feed Item** - Edit feed item content and properties
- **Contacts**
 - **Get Contact List** - View all contacts in the system
 - **Get Contact** - Fetch details of a specific contact
 - **Update Contact** - Edit contact information
- **Users**
 - **Get User List** - Retrieve list of Salesforce users
- **System**
 - **Describe SObject** - Get object metadata and field details for any Salesforce object type

Jira

Jira actions allow you to automate managing issues, projects, sprints and more.

- **Issues**
 - **Create Issue** - Create an issue in a project.
 - **Edit Issue** - Modify an existing issue in a project.
 - **Update Issue Status (Do Transition)** - Change task status of an issue.
 - **Delete Issue** - Delete an issue in a project.
 - **Get Issue** - View details of an issue in a project.
 - **Search Issues** - Searches for issues using JQL
 - **Search Statuses** - Search issue statuses.
 - **Download Issue Attachment** - View the contents of an attachment.
 - **Add Attachment** - Add an attachment to an issue.
 - **Get All Labels** - View all labels.
- **Comments**
 - **Add Comment** - Add new comment to an issue.
 - **Delete Comment** - Remove comment from an issue.

- **Update Comment** - Edit an existing comment on an issue.
- **Get Comments** - View issue comments.
- **Projects**
 - **Create Project** - Create new project in Jira.
 - **Update Project** - Modify project details.
 - **Delete Project** - Remove project from Jira.
 - **Get Project** - View project details.
 - **Get Issue Types For Project** - View project issue types.
 - **Search Projects** - Find visible projects.
- **Sprints**
 - **Create Sprint** - Create a sprint in a project.
 - **Update Sprint** - Update sprint details.
 - **Delete Sprint** - Delete a sprint in a project.
 - **Get Sprint** - View details of a sprint in a project.
 - **Move Issues to Sprint and Rank** - Assign an issue to a sprint.
 - **Move Issues to Backlog** - Move issues to backlog.
- **Users**
 - **Find Users** - Search for a Jira user.
 - **Find Assignable Users** - Returns a list of users that can be assigned to an issue.
 - **Get All Users** - List all Jira users.

ServiceNow

ServiceNow actions allow you to automate IT service management of incidents, change requests, and more.

- **Incidents**
 - **Create Incident** - Creates a new incident record.
 - **Read Incident** - Retrieves details of a specific incident.
 - **Update Incident** - Modifies an existing incident record.
 - **Delete Incident** - Removes a specific incident record.
- **Change Requests**

- **Create Change Request** - Creates a new change request record.
- **Read Change Request** - Retrieves details of a specific change request.
- **Update Change Request** - Modifies an existing change request record.
- **Delete Change Request** - Removes a specific change request record.
- **Knowledge Base Articles**
 - **Update Knowledge Base Article** - Modifies an existing knowledge base article.
- **Problem Records**
 - **Update Problem Record** - Modifies an existing problem record.
- **Attachments**
 - **Upload Binary Attachment** - Uploads a file as an attachment to a specified record.
 - **Upload Multipart Form Attachment** - Uploads a form attachment to a record.
 - **Retrieve Attachment Content** - Gets the binary content of a specific attachment.
 - **Retrieve Attachment Metadata** - Gets metadata for a specific attachment.
 - **Delete Attachment** - Removes an attachment from the system.
 - **Retrieve Metadata for Attachments** - Gets metadata for multiple attachments.
- **System**
 - **List Users** - Retrieves all existing records from system user table.
 - **List Choices** - Retrieves choice list values from the system.

SAP

SAP actions allow you to automate processes involving business partner data, material stocks, bills of material, product masters, and physical inventory.

- **SAP Business Partners**
 - **Get Business Partner** - Retrieves general data for business partners, providing a comprehensive view of partner information.
 - **Get Business Partner By ID** - Retrieves detailed business partner data using a specific business partner number as identifier.
 - **Get Business Partner Address** - Retrieves address details associated with business partners, including postal addresses and contact information.
 - **Get Business Partner Role** - Retrieves role information for business partners, showing their relationships and functions within the system.

- **Get Business Partner Role By ID** - Retrieves specific role data for a business partner using key identifier fields.
- **SAP Material Stock**
 - **Get Material Stock In Account Model Collection** - Gets material stock levels with details like batch, storage location, and special stock indicators.
- **SAP Bill Of Material**
 - **Get Material BOM Item** - Retrieves details of materials used in bills of materials including quantities, categories, and relationships.
- **SAP Product Master**
 - **Get Product Master Records** - View product master general details. Returns complete product master records with options to filter, sort and select specific data.
 - **Get Plant Data Of Product Master Record** - View plant details of the product. Access detailed plant-specific information for a given product including manufacturing, storage, and logistics data.
 - **Get Supply Planning Data By Product Number And Plant** - View supply planning details by product number and plant. Retrieves detailed planning parameters including lot sizes, safety stock levels, and procurement data.
- **SAP Physical Inventory**
 - **Get Phys Inventory Doc Item** - Retrieves physical inventory document items including counting results and inventory differences.

REST API Integration

The REST API Connection integration enables Quick Automate to interact with custom REST APIs and web services, allowing you to extend automation capabilities to external systems. This integration supports action execution for GET, POST, DELETE, PATCH and PUT HTTP methods.

Common Use Cases:

- Sending notifications to external systems
- Retrieving data from third-party services
- Updating records in external databases
- Triggering workflows in connected applications

Prerequisites

Before creating automations with REST API integration, ensure you have:

- **Configured REST API Integration**
 - Access Quick Suite console
 - Navigate to Integrations > REST API Connection
 - Configure your REST API endpoint with base URL
 - Complete authentication setup (see Authentication section below)
 - Full details: [REST API Connection integration](#)
- **Connected Integration to Automation Group**
 - Link your configured REST API integration to your automation group
 - Verify integration appears in available actions
 - Step-by-step guide: [AWS service connectors](#)
- **Authentication Configured**
 - OAuth 2.0 credentials obtained from your API provider

Authentication

For Quick Automate, this integration supports **OAuth 2.0** authentication. You must configure OAuth credentials before creating automations.

OAuth 2.0 Setup Steps:

- Obtain OAuth credentials (Client ID and Client Secret) from your identity provider
- In the Quick Suite console, navigate to your REST API integration settings
- Select OAuth 2.0 as the authentication method
- Enter your Client ID and Client Secret
- Configure the authorization URL and token URL provided by your API provider
- Complete the OAuth flow to authorize Quick Suite
- Test the connection to verify authentication

Important Notes:

- Quick Automate does not support bearer token authentication in request headers. Use OAuth configuration instead.
- Configure authentication at the integration level, not per automation action

Configuration

Setting Up Your REST API Integration

Step 1: Create Base URL Configuration

Step 2: Connect to Automation Group

Step 3: Access REST API Actions

- Open the Actions Panel in your automation
- Locate REST API action
- Drag and drop actions into your automation workflow

Note

REST API actions must be added through the Actions Panel; they are not available in the Build with Assistant interface

Properties Reference

- **Q Action Connector ID (Required)**
 - **Description:** Unique identifier that links the REST API action to your configured integration, including base URL and authentication settings.
 - **Format:** Auto-populated from your connected integration
 - **Example:** rest-api-connector-abc123
- **URI (Optional)**
 - **Description:** Additional path appended to the base URL to specify the exact API endpoint.
 - **Format:** Path string with or without leading slash (depends on the base url format)
 - **How It Works:**
 - Base URL: https://api.example.com
 - URI: /users/123

- Final URL: `https://api.example.com/users/123`

- **Examples:**

```
/users/123
prod/send-email
/v2/orders
customers/search
```

- **Additional Headers (Optional)**

- **Description:** Custom HTTP headers to include in the API request as key-value pairs.

- **Format:**

```
Key: Value
```

- **Common Examples:**

```
Content-Type: application/json
Accept: application/json
X-Custom-Header: custom-value
X-API-Version: 2.0
```

- **Important Limitations:**

- Do not include Authorization headers; authentication is handled via OAuth configuration
- Each header must be on a separate line
- Headers are case-sensitive

- **Request (Required)**

- **Description:** JSON-formatted data to send with POST or PUT requests.
- **Format:** Valid JSON object. {} for blank requests
- **POST Request Example (Creating a resource):**

```
{
  "title": "My New Post",
  "body": "This is the content of my new post.",
  "userId": 1,
  "tags": ["automation", "api"],
  "published": true
}
```

- **PUT Request Example (Updating a resource):**

```
{
  "status": "completed",
  "completedDate": "2026-01-27",
  "notes": "Task finished successfully"
}
```

- **Output Variable (Optional)**

- **Description:** Variable name to store the JSON response from the API call for use in subsequent automation actions.
- **Format:** Alphanumeric variable name (no spaces)
- **Example:** apiResponse, userData, orderDetails

Example of API usage

Example 1: GET Request - Retrieve Ticket Details

Scenario: Retrieve details of an existing support ticket from an external ticketing system

Configuration:

- **Q Action Connector ID:** rest-api-connector-tickets
- **HTTP Method:** GET
- **URI:** /tickets/ticketId
- **Additional Headers:**

```
Accept: application/json
```

- **Request Body:** {} (empty for GET requests)
- **Output Variable:** ticketDetails

Example 2: POST Request - Create New Record

Scenario: Submit a new support ticket to an external ticketing system

Configuration:

- **Q Action Connector ID:** rest-api-connector-tickets

- **HTTP Method:** POST
- **URI:** /tickets
- **Additional Headers:**

```
Content-Type: application/json
Accept: application/json
```

- **Request Body:**

```
{
  "title": "{{ticketTitle}}",
  "description": "{{ticketDescription}}",
  "priority": "{{ticketPriority}}",
  "requestedBy": "{{userEmail}}",
  "category": "technical_support"
}
```

- **Output Variable:** ticketResponse

Testing automations

Amazon Quick Automate provides comprehensive capabilities for testing your automations. This section describes the tools and features available for validating your automation behavior.

Test modes

Choose from two modes to test your automation:

- **Run mode** - Executes the automation from start to finish without stopping. Any breakpoints in the automation will be ignored. This mode is best for testing the complete flow of your automation.
- **Debug mode** - Executes the automation and pauses at any breakpoints you've set or if an exception occurs. This allows you to inspect the automation state, monitor variables, and validate behavior at specific points in your process.

Test pane

The Test pane includes includes two main tabs that help you monitor and control your automation test:

Setup tab

The Setup tab provides tools for managing your test configuration:

- Breakpoint manager - Lists all breakpoints currently set in your automation. Remove breakpoints directly from this view without returning to the canvas.
- Test navigation - Quick access links to view:
 - View runs - See all automation runs from testing
 - View cases - View cases created during testing
 - View tasks - Review any human-in-the-loop tasks generated

Monitor tab

The Monitor tab displays real-time information about your running automation:

- UI streaming - Watch browser automation actions execute in real-time
- Variable watcher - Monitor and modify variable values
- Logs - Track detailed automation progress

Test statuses

During testing, your automation will be in one of these statuses:

- Running - The automation is actively executing actions
- Paused (breakpoint) - Execution has stopped at a user-defined breakpoint
- Paused (exception) - Execution has stopped due to an error condition
- Completed - The automation has finished all steps successfully
- Failed - The automation encountered an unrecoverable error and stopped

Running and debugging

Test automations with inputs and outputs

When you test an automation that has input and output fields defined, Amazon Quick Automate validates the input values and displays both inputs and outputs as expandable cards in the **Monitor** pane alongside the execution logs.

Running an automation with inputs from the authoring studio

When you choose **Run** or **Debug** to test an automation with inputs and outputs, the authoring studio (Studio) displays a dynamic input form titled **Provide input values**. Amazon Quick Automate automatically generates this form based on the schema you defined.

The form includes:

- Field names from your schema.
- Input controls appropriate for each data type (text fields for strings, number inputs for integers, and file upload buttons for files).
- Required field indicators (marked with an asterisk).
- Default values pre-populated where you defined them.
- Client-side validation for data types, required fields, and file size limits (maximum 5 MB for files).

To run an automation with inputs:

- Open your automation in Studio and choose **Run** or **Debug**.
- In the **Provide input values** form, enter values for each field:
 - For text and number fields, type the values directly
 - For file fields, choose **Browse** to select a file from your local machine
 - Optional fields can be left blank unless you want to override default values
- Choose **Start** to begin the automation with your provided inputs.

Amazon Quick Automate performs two layers of validation:

- **Client-side validation** – Checks for blank required fields, correct data types, and file size limits before execution starts
- **Runtime validation** – Validates the complete input payload against your JSON schema during execution

Viewing inputs and outputs

After the automation runs, input and output values appear as collapsible cards in the **Monitor** pane:

- **Input card** – Displays at the beginning of the run logs, showing all values you provided when starting the automation
- **Output card** – Displays at the end of the run logs after execution completes, showing all values your automation produced

Both cards are collapsed by default. When you expand them, you can:

- View the complete schema definition and actual values
- Download any file objects included in the inputs or outputs
- Copy the structured data as JSON
- Download the complete payload as a JSON file

Note

If the automation terminates with an error, the output card won't be available.

Setting breakpoints

Breakpoints allow you to pause your automation at specific points:

To add a breakpoint:

- Hover over an action or step on the canvas
- Click **Add breakpoint**
- The breakpoint icon appears on the action or step

To remove a breakpoint:

- Click **Remove breakpoint** on the action
- Or use the breakpoint manager in the Setup tab

Note

Set up breakpoints before starting debug mode.

Debug controls

When your automation is paused, you have several control options:

- Next - Run the next action in your automation and pause again
- Continue - Resume the run until the next breakpoint is reached
- Stop - End the current test run
- Ignore - If paused on an exception, this will clear the exception and continue as if that action was successful. This is useful if the exception is not blocking the rest of your test and you want to continue.
- Retry - If paused on an exception, this will attempt to run the current action again. This is useful if the issue is temporary or if you want to change the value of a variable and retry with the new value

Deactivating steps

To temporarily disable any actions or steps without removing them, you can deactivate them before starting your test. Deactivated actions or steps will be skipped during testing.

- Select the step to deactivate
- Click the **Menu options**
- Choose **Deactivate**
- The step appears grayed out

To reactivate:

- Select the deactivated step
- Click the **Menu options**
- Choose **Activate**

Variable watcher

The Variable watcher in the Monitor tab provides visibility to inspect data being used by your automation during the test. To inspect a variable:

- Open the Variable watcher in the Monitor tab while the test is paused

- Click **Add** to search and select variables to inspect
- View the current value for the variable
- Modify values to test different scenarios:
 - Click the **pencil icon** next to any variable
 - Enter a new value to test
 - Continue the run to proceed with the new value

This feature is particularly useful for:

- Testing different data conditions
- Validating error handling
- Simulating edge cases

Troubleshooting

Fix with Assistant

When your automation encounters an exception, the Automation Assistant can help identify and resolve issues:

- Click **Fix with Assistant** when paused at an error
- The Assistant will:
 - Analyze the error message
 - Review the run logs
 - Check its troubleshooting knowledge base
 - Provide a clear explanation of the issue
 - Suggest potential solutions
- Chat with the Assistant to:
 - Ask questions about the error
 - Update the automation based on any of the recommended changes

Test metrics

After your test completes, Amazon Quick Automate provides comprehensive test metrics to help understand automation performance and validate test coverage. View key metrics about your test run:

- **Total running time** - How long the test took from start to finish
- **Number of completed cases** - Count of cases that finished processing (excluding those created or pending)
- **Average time per case** - The processing average time for each completed case
- **Case success rate** - Percentage of cases that completed without exceptions
- **Count of tasks created** - Number of human-in-the-loop tasks generated during the test
- **Action coverage** - Percentage of actions in your automation that executed during the test, helping identify untested paths

Deploying automations

After you create and test your automations in Amazon Quick Automate, the next step is to **deploy** them so they can run regularly on a trigger. Deployment makes the automation operational and ready for execution. You can add a trigger to a deployed automation to run it at a predefined schedule. Deployment involves configuring runtime settings, assigning users for human-in-the-loop tasks, verifying credentials and integrations.

This section explains each step of the deployment process and how to configure your automation for reliability, security, and optimal performance.

Prerequisites

Before deploying your automation, ensure that the following steps are complete:

- **Automation is tested thoroughly** - Validate your automation through end-to-end testing to confirm that all logic, actions, and agent interactions work as expected.
- **Version committed for deployment** - Only committed automation versions can be deployed. Review your changes, finalize the version, and commit it before proceeding.
- **Integrations configured** - If your automation interacts with external applications such as Salesforce or Jira via APIs, make sure all necessary integrations are configured.
 - Navigate to **Connections** → **Integrations** in the left panel to create new integrations.

- Currently, only integrations available under the **Actions** tab are supported in Amazon Quick Automate.
- Once an integration action is created, associate it with the **Automation Group** where it will be used.
- The associated actions will then appear in the canvas. During deployment, you can select the appropriate connection to be used by the deployed automation.
- **Credentials configured** - Verify that all credentials required by your automation are correctly set up.

Deploying an automation

You can deploy an automation directly from the Canvas by clicking Deploy, or by navigating to the Deployment tab on the automation's landing page. Once you initiate deployment, the system guides you through a series of steps to complete the configuration and release process.

Release details

On the **Release Details** page, select the version of the automation you want to deploy. Only **committed versions** are available for deployment and will appear in the dropdown list.

Additional settings

Additional settings include:

- **Runtime configuration**
- **Tasks (for HITL assignments)**
- **Access**

Runtime configuration

Runtime configurations are parameters that may differ between environments such as development, testing, and production.

For example, an automation step that sends an email may use your personal email address during testing but should switch to a shared team address in production. Such environment-dependent values can be defined as **runtime configurations** when authoring the automation.

At deployment time, you can review and override these configurations to ensure the automation runs correctly in the intended environment. Runtime parameters can include:

- Email addresses or notification recipients
- File paths or URLs specific to the environment

This flexibility helps maintain a single automation definition across environments while adapting key parameters as needed.

Tasks

Select the resolver users or user groups for each of the tasks that require human-in-the-loop.

Access (Credentials and Connections)

Automations often need to connect to external systems, databases, or services. Amazon Quick Automate provides secure methods for managing credentials and connections without embedding sensitive information within your automation logic.

Connection and credential data are securely stored and encrypted, and made available to workflows at runtime without exposing them to authors or end users. This design ensures strong separation of secrets from automation definitions, improving both security and maintainability.

You can store and use two primary types of credentials:

- **Website credentials** - Used for UI automation steps that require website logins (username and password).
- **Action credentials (Integrations)** - Used for connecting AWS services (e.g. S3) or external systems through configured integrations (e.g., Salesforce, Jira).

At deployment, ensure that the correct credentials and connections are selected so the automation can access all required systems securely and successfully.

Setting up triggers

Triggers determine when and how your automations run. You can configure automations to start based on a predefined schedule or invoke them programmatically through the Amazon QuickSight API.

Schedules

You can configure automations to start based on a predefined schedule. To setup a trigger:

- On the Deployment page, click "Create Trigger" and configure the rules.
- Select the frequency
- Select the start date and time (Note that the actual execution will initiate within 15 minutes of the start time selected)
- Select the end date and time
- Select the timezone
- Amazon Quick Automate provided built-in scalability. Select the number of parallel executions of the automation (you can select a maximum of 10 parallel executions per trigger and 50 across all automations within an account. Please reach out to AWS)
- For complex scheduling needs, you can use cron expressions to define precise running patterns. For example, to run an automation at 2:30 AM every Monday, Wednesday, and Friday, you would use the cron expression: `30 2 * * 1,3,5`.

API triggers

The Automation Job APIs enable you to programmatically start and monitor automation jobs from external applications. You can use these APIs to invoke deployed automations with custom input payloads and retrieve execution results through the AWS SDK and AWS CLI.

The Automation Job APIs include two operations:

- **StartAutomationJob** – Starts a new job for a deployed automation with an optional input payload.
- **DescribeAutomationJob** – Retrieves the status, timestamps, and optional input and output payloads for a specified job.

These APIs are part of the Amazon QuickSight service namespace. You access them through the `quicksight` namespace in the AWS SDK and AWS CLI.

For more information about AWS SDKs and toolkits, see the [AWS Getting Started Resource Center](#).

Prerequisites

Before you call the Automation Job APIs, complete the following setup steps.

Find your automation identifiers

To call the Automation Job APIs, you need the following identifiers:

- **AWS account ID** – Your 12-digit AWS account ID.
- **Automation group ID** – The unique ID of the automation group that contains your automation.
- **Automation ID** – The unique ID of the automation you want to invoke.

You can find the automation group ID and automation ID in the Deployments section when you open any automation and have a deployed version.

To find your identifiers

- Sign in to Amazon Quick Automate.
- In the left navigation pane, choose **Automations**.
- Choose the automation group that contains your automation.
- Choose the automation name to open the automation editor.
- Choose the **Deployments** tab.
- Choose the **Actions** (:) and **View deployment details** to get the **Automation ID** and **Group ID** at the top just below the heading Deployment Details.

Configure IAM permissions

The Automation Job APIs require IAM permissions attached to your IAM identity (user, role, or group). Each operation requires a separate permission.

Permission for StartAutomationJob

Attach a policy that grants the `quicksight:StartAutomationJob` action. Scope the resource ARN to the specific automation that you want to allow.

```
{
  "Version": "2012-10-17"
  ,
  "Statement": [
    {
      "Sid": "AllowStartAutomationJob",
      "Effect": "Allow",
      "Action": [
        "quicksight:StartAutomationJob"
      ],
      "Resource": [
```

```

        "arn:aws:quicksight:us-west-2:111122223333:automation-group/EXAMPLE-
        GROUP-ID/automation/EXAMPLE-AUTOMATION-ID"
    ]
}
]
}

```

Permission for DescribeAutomationJob

Attach a policy that grants the `quicksight:DescribeAutomationJob` action. The resource ARN for this operation includes the job ID segment. To allow describing any job under an automation, use a wildcard (*) for the job segment.

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowDescribeAutomationJob",
      "Effect": "Allow",
      "Action": [
        "quicksight:DescribeAutomationJob"
      ],
      "Resource": [
        "arn:aws:quicksight:us-west-2:111122223333:automation-group/EXAMPLE-
        GROUP-ID/automation/EXAMPLE-AUTOMATION-ID"
      ]
    }
  ]
}

```

AWS recommends scoping each permission to the most specific resource ARN possible to follow the principle of least privilege. For more information, see the following:

- [IAM policy examples for Amazon QuickSight](#)
- [Actions, resources, and condition keys](#)
- [IAM JSON policy elements](#)

Deploy your automation

The `StartAutomationJob` API targets deployed automations only. You must commit and deploy your automation in the Amazon Quick Automate console before you can invoke it through the API.

StartAutomationJob

Starts a new job for a deployed automation. The operation is asynchronous — Amazon Quick Automate accepts the request, queues the job, and returns a **JobId** immediately. The automation runs in the background. If the automation has an input schema defined, Amazon Quick Automate validates the **InputPayload** against the schema before accepting the job. An invalid payload results in an **InvalidParameterValueException**.

Request syntax

```
POST /accounts/AwsAccountId/automation-groups/AutomationGroupId/
automations/AutomationId/jobs HTTP/1.1
Content-type: application/json

{
  "InputPayload": "string"
}
```

Request parameters

- **AwsAccountId** (String, required) – Your AWS account ID (12 digits).
- **AutomationGroupId** (String, required) – The ID of the automation group (UUID).
- **AutomationId** (String, required) – The ID of the automation to run (UUID).
- **InputPayload** (String, optional) – Input for the job as a JSON string.

Response elements

- **Arn** (String) – ARN of the automation job.
- **JobId** (String) – ID of the started job. Use this with `DescribeAutomationJob` to track status.
- **Status** (Integer) – HTTP status code of the response.
- **RequestId** (String) – AWS request ID.

Errors

- **AccessDeniedException** – Insufficient permissions or invalid credentials.
- **InvalidParameterValueException** – One or more parameters has a value that isn't valid.
- **ResourceNotFoundException** – Automation group or automation not found.

- **LimitExceededException** – A limit is exceeded.
- **ThrottlingException** – Request was throttled.
- **InternalFailureException** – Internal service error.

For more details, see the [API reference guide for StartAutomationJob](#).

DescribeAutomationJob

Retrieves the status and details of a specified automation job, including execution timestamps and optional input and output payloads. Use this operation to poll for job completion after calling **StartAutomationJob**.

The response does not include input and output payloads by default. Set the **IncludeInputPayload** and **IncludeOutputPayload** query parameters to `true` to request them. Amazon Quick Automate returns the output payload only when all of the following conditions are met:

- **IncludeOutputPayload** is `true`.
- **JobStatus** is `SUCCEEDED`.
- The automation produced output values.

Request syntax

```
GET /accounts/AwsAccountId/automation-groups/AutomationGroupId/
automations/AutomationId/jobs/JobId?
IncludeInputPayload=boolean&IncludeOutputPayload=boolean HTTP/1.1
```

Request parameters

- **AwsAccountId** (String, required) – Your AWS account ID (12 digits).
- **AutomationGroupId** (String, required) – The ID of the automation group (UUID).
- **AutomationId** (String, required) – The ID of the automation (UUID).
- **JobId** (String, required) – The ID of the job returned by **StartAutomationJob** (UUID).
- **IncludeInputPayload** (Boolean, optional) – Include the input payload in the response. Default: `false`.
- **IncludeOutputPayload** (Boolean, optional) – Include the output payload in the response. Default: `false`.

Response fields

- **Arn** (String) – ARN of the automation job.
- **CreatedAt** (Timestamp) – When the job was created (epoch seconds).
- **StartedAt** (Timestamp) – When the job started running (epoch seconds).
- **EndedAt** (Timestamp) – When the job finished (epoch seconds).
- **JobStatus** (String) – QUEUED, RUNNING, SUCCEEDED, FAILED, or STOPPED.
- **InputPayload** (String) – The input payload. Omitted from response unless `IncludeInputPayload` is true.
- **OutputPayload** (String) – The output payload. Omitted from response unless `IncludeOutputPayload` is true.
- **RequestId** (String) – AWS request ID.

Errors

- **AccessDeniedException** – Insufficient permissions or invalid credentials.
- **InvalidParameterValueException** – One or more parameters has a value that isn't valid.
- **ResourceNotFoundException** – Job, automation, or automation group not found.
- **ThrottlingException** – Request was throttled.
- **InternalFailureException** – Internal service error.

For more details, see the [API reference guide for DescribeAutomationJob](#).

Run automations with API triggers using input and output

You can include input values as a JSON payload when you start an automation job through the API. Amazon Quick Automate validates the payload against the automation's input schema before execution begins, and returns structured output values when the job completes.

Prerequisites

Before you run an automation with API triggers, make sure the following are in place:

- The automation is deployed.
- The automation has an input schema defined in the Start node.
- You have the automation group ID and automation ID.

Sending input values through the API

To include input values, pass a JSON payload in the **StartAutomationJob** request using the **InputPayload** parameter.

Sample Request

```
POST /accounts/123456789012/automation-groups/a1b2c3d4-e5f6-7890-abcd-ef1234567890/
automations/11111111-2222-3333-4444-555555555555/jobs
{
  "InputPayload": "{\"customer_id\":\"C-98765\",\"threshold\":100,\"region\":\"us-
east-1\"}"
}
```

Sample Response

```
{
  "Arn": "arn:aws:quicksight:us-west-2:123456789012:automation-group/a1b2c3d4-
e5f6-7890-abcd-ef1234567890/automation/11111111-2222-3333-4444-555555555555/job/
aaaaaaaa-bbbb-cccc-dddd-eeeeeeeeeeee",
  "JobId": "aaaaaaaa-bbbb-cccc-dddd-eeeeeeeeeeee",
  "Status": 200,
  "RequestId": "req-12345678-abcd-efgh-ijkl-123456789012"
}
```

Note

The **InputPayload** value must be a JSON-serialized string that contains the input field values matching your automation's input schema.

Retrieving output values through the API

You can retrieve output values by calling the **DescribeAutomationJob** API. To include input and output payloads in the response, set the **includeInputPayload** and **includeOutputPayload** query parameters to **true**.

The response includes the **OutputPayload** field only when the automation completes successfully.

Sample Request

```
GET /accounts/123456789012/automation-groups/a1b2c3d4-e5f6-7890-abcd-ef1234567890/
automations/11111111-2222-3333-4444-555555555555/jobs/aaaaaaaa-bbbb-cccc-dddd-
eeeeeeeeeeee?includeInputPayload=true&includeOutputPayload=true
```

Sample Response

```
{
  "Arn": "arn:aws:quicksight:us-west-2:123456789012:automation-group/a1b2c3d4-
e5f6-7890-abcd-ef1234567890/automation/11111111-2222-3333-4444-555555555555/job/
aaaaaaaa-bbbb-cccc-dddd-eeeeeeeeeeee",
  "JobStatus": "SUCCEEDED",
  "CreatedAt": "2026-03-11T09:59:50Z",
  "StartedAt": "2026-03-11T10:00:00Z",
  "EndedAt": "2026-03-11T10:05:00Z",
  "InputPayload": "{\"customer_id\":\"C-98765\",\"threshold\":100,\"region\":\"us-
east-1\"}",
  "OutputPayload": "{\"result\":\"success\",\"records_processed\":42,\"summary\":
{\"passed\":40,\"failed\":2}}",
  "RequestId": "req-12345678-abcd-efgh-ijkl-123456789012"
}
```

Note

The **OutputPayload** field is available only after the automation completes successfully (when **JobStatus** is SUCCEEDED). If the automation fails or is still running, the response does not include the **OutputPayload** field.

Deploy and run automations with inputs and outputs

This section describes how to deploy automations that use input and output schemas, run them manually or on a schedule, and view run results.

View schema information before deployment

Before deploying an automation, you must commit it to create a version. For more information, see [maintaining automation versions](#).

After you commit an automation with input or output schemas, you can view the schema details in the **Data model** section of the version details page. This section displays the complete field definitions for both inputs and outputs.

Deploy automations with input and output schemas

When you deploy an automation that has input or output schemas defined, the deployment dialog box displays schema information in the **Data model** section, including:

- Input and output field definitions (name, data type, description, required or optional status, and default values).
- Copyable JSON schemas for inputs and outputs that API callers can use for integration.

Important

The schema is frozen at deployment time along with the automation code. If you update the schema after deployment, you must redeploy the automation for the changes to take effect.

To complete the deployment, review the schema information and choose **Deploy**.

Run deployed automations manually

To start a deployed automation with inputs, complete the following steps:

- On the deployment page, choose the **Actions** menu (:).
- Choose **Run now**.
- In the **Provide input values** dialog box, enter values for each input field.
- Choose **Start** to start the automation run.

The input form is identical to the one you use when you test in Studio.

Run automations with scheduled triggers

For scheduled triggers, you provide input values when you create or edit the trigger. The stored values are passed to the automation each time the trigger fires.

You can do the following:

- Create multiple triggers with different input values for the same automation
- Edit input values for a trigger without redeploying the automation

To configure inputs for a scheduled trigger, complete the following steps:

- Navigate to the deployment page for your automation.
- Choose **Create Trigger**.
- In **Step 1: Set Trigger**, configure the schedule.
- In **Step 2: Define Run**, enter values for each input field in the auto-generated form.
- Choose **Next**, review your configuration, and choose **Create Trigger**.

 Note

If you create a trigger with a specific schema version and later deploy the automation with an updated schema, you must modify the trigger. The trigger becomes incompatible when the schema version changes and the automation runs will fail.

Viewing inputs and outputs for deployed runs

After an automation run completes (whether started manually or by a scheduled trigger), input and output values are available as structured artifacts in the logs panel on the **Runs** page.

The artifacts appear as dedicated cards:

- **Input artifact** – Displays at the top of the logs panel
- **Output artifact** – Displays at the bottom of the logs panel (only available if the automation completes successfully)

Both cards are collapsed by default. When you expand them, you can:

- View the complete schema and actual values
- Download any file objects included in the inputs or outputs
- Copy the structured data as JSON
- Download the complete payload as a JSON file

Export and import automations

You can use the import and export feature in Amazon Quick Automate to share and migrate automation versions across automation groups, AWS accounts, and AWS Regions. Use this feature to promote automations from development to production environments, create backups, or share automations with other teams.

When you export an automation version, Amazon Quick Automate packages it into a secure, encrypted link that you can use to transfer the automation between environments. The export bundle contains the following:

- The automation workflow
- Runtime configuration settings
- Process step descriptions and metadata

Note

Dependencies such as connectors, credentials, and task queue configurations are not included in the export bundle. You must configure these separately in the destination environment.

Use cases

The import and export feature supports the following use cases:

- Move tested automations from development to production accounts.
- Deploy automations across different AWS Regions.
- Create point-in-time snapshots of automation versions for disaster recovery.
- Share automation versions across different automation groups within the same account.

Prerequisites

Before you import or export automations, make sure that you have the following:

- Access to Amazon Quick Automate.

- Appropriate permissions to create and manage automations in both the source and destination environments.
- Amazon Quick Automate enabled in both the source and destination AWS accounts (for cross-account transfers).

Export an automation

To export an automation version, complete the following steps.

- Navigate to your automation in **Amazon Quick Automate**.
- Choose the actions menu (:) next to **Edit automation**.
- Choose **Export version**.
- Choose the version from the dropdown list and wait for the link to be generated.
- Choose **Copy link** or copy the link manually.

Amazon Quick Automate generates a secure, pre-signed download link for your automation. Amazon Quick Automate encrypts the bundle by using AWS KMS with envelope encryption to protect your automation content.

Export link validity

Export links are valid for 12 hours from the time of creation. After this period, the link expires and you must generate a new export link. You can reuse the same export link multiple times within the 12-hour window without creating a new export.

Import an automation

To import an automation version, complete the following steps.

- Navigate to the destination automation group and create a blank automation.
- Choose the actions menu (:) next to **Edit automation**.
- Choose **Import version**.
- Paste the exported link (the pre-signed URL that you received during export).
- Choose **Start**.

Considerations

- You can export one automation version per operation.
- Connectors, credentials, and other dependencies are not included in the export. You must configure these separately in the destination environment.
- For cross-account transfers, both the source and destination accounts must have Amazon Quick Automate enabled.

Monitoring performance

Amazon Quick Automate provides comprehensive monitoring capabilities that help you track performance, audit and troubleshoot your automations. This section describes how to use the performance monitoring features in Amazon Quick Automate.

Prerequisites

Before using the Performance dashboard, you need:

- **Active automations** - At least one automation must be run to see metrics
- **Automation group access** - Permission to view the automation groups you want to monitor

Automation summary dashboard

The Automations tab provides high-level metrics and key performance indicators across your automation portfolio.

Automation summary charts

The dashboard includes three main performance visualizations:

- **Automation Distribution** - Shows the number of deployed automations organized by automation group, helping you understand your automation footprint across the organization
- **Success Rate** - Displays the percentage of cases that completed successfully versus those that encountered exceptions, providing insight into automation reliability
- **Time Savings** - Calculates total hours saved based on successful case completions multiplied by the estimated time savings configured for each automation project

Automation details table

Each automation is listed with key metrics and status information:

- **Automation name** - Name and version of the deployed automation
- **Latest run** - Most recent execution status and timestamp
- **Case metrics** - Count of cases processed, broken down by:
 - Successfully completed cases
 - Cases with business exceptions
 - Cases with system exceptions
- **Tasks** - Number of human-in-the-loop tasks created by the automation

Runs dashboard

The Runs tab provides detailed visibility into individual automation executions and their outcomes.

Run statuses

Automations can be in one of these run statuses:

- **Waiting** - The automation is queued and awaiting available system resources to begin execution. This is typically a brief transitional state.
- **In Progress** - The automation is actively executing its configured steps. You can monitor real-time progress through the logs.
- **Completed** - The automation has finished all steps successfully.
- **Failed** - The automation encountered an error it could not recover from and stopped execution. Detailed error information is available in the logs.
- **Stopped** - A user manually interrupted the automation execution using either the Exit or Terminate action.

Note

Count of incomplete runs includes Waiting and In Progress.

Run details table

Each run provides comprehensive execution information through the following columns:

- **Status** - Current run status (as defined in Run Status Types)
- **Automation** - Name of the automation being run
- **Version** - Specific version number of the automation which ran
- **Start time** - When the automation run began
- **End time** - When the automation completed or stopped
- **Duration** - Total running time of the automation
- **Case metrics** - Count of cases processed, broken down by:
 - Successfully completed cases
 - Cases with business exceptions
 - Cases with system exceptions
- **Tasks** - Number of human-in-the-loop tasks created in the run

Note

Data is refreshed when first navigating to the page. Click the refresh button to pull the latest data at any point.

Cases dashboard

The Cases tab enables tracking of individual case records through their lifecycle.

Case charts

The dashboard includes two main performance visualizations:

- **Total cases created** - Shows the number of cases created over time, helping you understand your automation volumes.
- **Top exceptions** - Displays the top 5 exception reasons and the relative count of each, helping you prioritize your optimization efforts to improve success rates.

Case statuses

Cases can be in one of these processing statuses:

- **Ready** - The case has been created and is waiting to be picked up for processing. This is the initial state of all new cases.
- **In Progress** - The case is actively being processed by the automation. You can monitor real-time progress through the logs.
- **Pending Resolution** - The case is waiting for a human-in-the-loop task to be completed before processing can continue. Cases move back to Ready status after tasks are resolved.
- **Successful** - The case has completed without any exceptions.
- **Failed** - The case encountered one of the following exception types:
 - **Business Exception** - The case encountered a handled business rule violation and stopped processing. Detailed exception information is available in the logs.
 - **System Exception** - The case experienced a technical error and stopped processing. Detailed error information is available in the logs.

For detailed information about case handling and orchestration patterns, see [Orchestration actions](#).

Search and filtering

The Performance dashboard includes search and filtering capabilities available across all monitoring views to help you find specific information. Filter by:

- Time range
- Status
- Automation group
- Automation

Note

You can also search on the cases page by reference name or exception reason.

Environment selection

Switch between viewing metrics from:

- Test - Data from automation testing in development environment
- Deployed - Data from deployed automations

Note

Your selected environment applies across all dashboard tabs until changed.

CloudWatch

Amazon CloudWatch provides monitoring and observability capabilities for your automations in Amazon Quick Automate. This section explains the metrics available in CloudWatch, how to view the metrics, and how to set up alarms for these metrics.

Prerequisites

IAM roles and permissions

To grant CloudWatch access following least-privilege principles:

- Create an IAM role or group, preferably using AWS IAM Identity Center.
- Attach the AWS managed policy CloudWatchFullAccess for full monitoring and alarm capabilities.
- Optionally, add CloudWatchLogsReadOnlyAccess for log querying without deletion permissions.
- Assign users to this role or group to enable CloudWatch operations without broader administrative access.

For detailed step-by-step instructions, see [Getting setup in CloudWatch](#) and [CloudWatch permissions guide](#).

Accessing CloudWatch

- Navigate to the AWS Management Console.
- Enter CloudWatch in the services search bar.

- Choose CloudWatch from the results to open the CloudWatch console.

Accessing CloudWatch metrics for Amazon Quick Automate

- Open the CloudWatch console.
- In the left navigation pane, choose **Metrics, All metrics**.
- Choose the service namespace **QuickSight**.
- Choose the relevant dimensions, **AutomationGroupId** or **AutomationId**, to view the metrics.
- Use the search bar to search or filter by metric name or dimension.

For more detailed information, see the [CloudWatch metrics documentation](#).

Available metrics

Amazon Quick Automate publishes metrics to CloudWatch under the **QuickSight** namespace. Metrics are grouped by either **AutomationId** or **AutomationGroupId** dimensions.

Note

CloudWatch is a Regional service. To view metrics, ensure that you are in the same AWS Region where your automations run.

AutomationId metrics

These metrics track individual automation performance:

- **FailedAutomationRunCount** - The total number of automation runs that failed during execution. Use this metric to identify problematic automations that require attention or debugging.
- **SuccessfulAutomationRunCount** - The total number of automation runs that completed successfully. Use this metric to track automation reliability and success rates.
- **AutomationRunDuration** - The time taken for each automation run to complete. Use this metric to identify performance bottlenecks and optimize automation efficiency.
- **StoppedAutomationRunCount** - The number of automation runs that were manually stopped or terminated before completion. This can indicate user intervention or system-initiated stops.

- **TotalAutomationRunCount** - The aggregate count of all automation runs, regardless of status. Use this metric to monitor overall automation activity and usage patterns.

AutomationGroupId metrics

These metrics provide a broader view across all cases:

- **TotalCaseCount** - The total number of cases processed by the automation system. Use this metric to track overall workload volume.
- **CompletedCaseCount** - The number of cases that were successfully processed and completed. Use this metric to measure throughput and productivity.
- **FailedCaseCount** - The number of cases that failed during processing. Use this metric to identify systemic issues or patterns in failures.
- **CaseDuration** - The time taken to process cases from start to completion. Use this metric to assess overall system performance and identify opportunities for optimization.

Finding AutomationId and AutomationGroupId

To find the AutomationId and AutomationGroupId for your automation:

- Sign in to Amazon Quick Automate.
- In the left navigation pane, choose **Automations**.
- Choose the automation group that contains your automation.
- Choose the automation name to open the automation.
- Choose the **Deployments** tab.
- Choose the **Actions** (:) and **View deployment details** to get the **Automation ID** and **Group ID** at the top just below the heading Deployment Details.

Setting up CloudWatch alarms

Use these metrics to monitor the health of your automations or set up alarms in CloudWatch. For example, you can create alarms for the following conditions:

- More than five automation runs fail within a specified period.
- More than 10 cases fail within a specified period.
- Case duration exceeds 30 minutes.

For more information about creating alarms, see [CloudWatch alarms](#).

Human-in-the-loop task center

Human-in-the-loop (HITL) tasks enable you to incorporate human judgment at critical points in your automated processes. The Task Center provides a centralized interface for managing HITL tasks. Users can review, resolve, and track tasks that require human intervention.

Overview

The Task center displays tasks requiring human attention and provides metrics on task processing.

Task metrics

The dashboard shows key performance indicators:

- **Open tasks** - Total number of tasks awaiting completion
- **Critical tasks** - Count of high-priority tasks requiring attention
- **Overdue tasks** - Number of tasks past their due date

Task details table

Each task provides comprehensive information through the following columns:

- **Title** - Brief description of what needs to be done
- **Status** - Current task state (Open, Completed, Archived)
- **Due date** - When the task should be completed (includes overdue warning)
- **Severity** - Task priority (Low, Moderate, Critical)
- **Assignee** - User or group responsible for the task
- **Creation time** - When the task was created
- **Completion time** - When the task resolution was submitted by a user

Prerequisites

Before working with tasks, you need:

- Task permissions - Access granted through automation deployment settings
- Active automations - At least one automation with HITL tasks must be run

Types of tasks

Chat tasks

Interactive conversations where AI agents gather input through natural language. To participate in a chat task:

- Select the task
- Review the conversation history
- Respond to agent questions in the chat interface
- The agent will close the task once it has sufficient information

Note

Chat tasks are closed automatically by the AI agent. You can archive unresolvable tasks, which the agent will treat as exceptions.

Form tasks

Tasks requiring structured input through form fields. To complete a form task:

- Select the task
- Fill in the required form fields
- Click Submit when finished

Note

Enable "Automatically open next task" to streamline processing multiple tasks. Use "Save as Draft" to preserve partial progress.

Managing tasks

Task statuses

Tasks can be in one of these states:

- Open - Awaiting human input
- Completed - Successfully submitted
- Archived - Removed from active view

Task actions

Manage tasks through these available actions:

- Submit - Complete a form task
- Save as Draft - Preserve partial form progress
- Archive - Remove task from active view
- Unarchive - Restore archived task

Note

Only Open tasks can be archived.

Task assignment

Distribute work by assigning tasks:

- Select a task
- Click the assignee field
- Search for and select users or groups
- Click Remove assignee to clear assignment

Note

Users must have appropriate permissions to be assigned tasks.

Related information

Access connected resources through the Sources menu:

- Case - View related case details (if applicable)
- Automation - See the automation that created the task

Note

Source access requires appropriate permissions.

Search and filtering

Locate specific tasks using:

- Search - Find tasks by title
- Filters:
 - Status
 - Due date
 - Severity
 - Assignee
 - Creation date

Note

Click refresh to update task data at any time.

Environment selection

Switch between viewing tasks from:

- Test - Tasks from automation testing
- Deployed - Tasks from deployed automations

Build web applications with apps in Amazon Quick

You can use apps in Amazon Quick to build interactive web applications by describing what you need in natural language. You tell the apps in Quick agent what your app should do, who it is for, and what data it needs. The agent builds a working app in real time while you watch. You refine it through ongoing conversation. You can then publish and share the app for daily use across your organization.

With Apps in Quick you can bridge the gap between data visualization and custom application development. Whether you need a dashboard, an internal tool, a data entry interface, or a document viewer, you can create it conversationally and publish it for your team in minutes.

Key capabilities

With Apps in Quick, you get the following capabilities:

- **Conversational authoring** — Describe what you want, and the AI builds it.
- **Live dashboard embeds** — Embed interactive Amazon Quick Sight visuals directly in your app.
- **Connectors** — Call external APIs and services from your app. For more information, see [Work with integrations in Amazon Quick](#).
- **Built-in AI inference** — Add AI-powered features such as text summarization, classification, and content generation using foundation models.
- **Spaces integration** — Read and manage documents and resources from Amazon Quick spaces. For more information, see [Organize, collaborate, and share resources with spaces in Amazon Quick](#).
- **Persistent data storage** — Save and retrieve data across sessions with built-in key-value storage.
- **Enterprise authentication** — Inherits Quick's security model with single sign-on support.
- **One-click publishing** — Share apps with your team instantly.

How apps in Quick works

Building an app follows this workflow:

1. **Describe** — Tell the AI what you want to build.
2. **Generate** — The AI writes the code in real time.

3. **Preview** — See a live preview instantly and test the app.
4. **Iterate** — Refine with follow-up requests.
5. **Publish** — Share with your team in one click.

Getting started with apps in Quick

You can create your first app in apps in Amazon Quick in minutes. This section walks you through each step.

Prerequisites

Before you begin, confirm the following:

- You have an active Amazon Quick account.
- Your subscription includes Author, Professional, Author Pro, Enterprise, or Admin Pro access.

Open the apps in Quick landing page

To open the apps in Quick landing page:

1. Sign in to Amazon Quick.
2. In the navigation pane, choose **Apps**.

The apps in Quick landing page provides a prompt area where you describe the app you want to build. Rotating example prompts suggest ideas you can use or customize.

Below the prompt area, you can browse your existing apps in two views: **Library** (all apps you have access to, including apps shared with you) and **My apps** (apps you created). You can search for apps by name or owner and sort results by last visited, last updated, or name.

Create an app

To create an app:

1. In the prompt area, enter a description of the app you want to build. Alternatively, choose one of the example prompts to start with a pre-written description, then edit it before proceeding.
2. Choose **Generate**.

3. Review the live preview in the apps in Quick editor as the agent builds your app. Use the chat panel to provide follow-up instructions.

Write your description as if you were explaining the app to a colleague. Include the purpose, the audience, and the data it needs. A detailed description produces a better first version and reduces iteration.

Edit an app

The apps in Quick editor is a two-panel workspace. The preview panel shows a live, interactive rendering of your app. The chat panel is where you converse with the apps in Quick agent to build, modify, and debug the app.

To edit an app:

1. Enter an instruction in the chat panel.
2. Wait for the agent to update the preview with your changes.
3. Interact with the preview to test buttons, forms, navigation, and other features.

The chat conversation persists across sessions. If you close the editor and return later, your previous messages are still there. Choose **Load older messages** to see the full history.

Manage versions

Apps in Quick automatically creates a new version each time the agent changes your app. You do not need to save manually. A version selector in the editor lets you browse all previous versions and preview any of them.

To restore a previous version:

1. Select the version you want to review from the version selector.
2. Confirm the preview shows the state you want.
3. Choose **Restore to this version**.

Restoring a version does not affect the published version. You must publish again for end users to see the change. For more information about recovering from a broken app, see [Recovering from a broken app](#).

Review app settings

The app settings panel provides two views:

- **Overview** — Shows the app name (editable), a unique app ID, creation date, last published date, the published URL, the public URL (if published publicly), and an editable description. You can also delete the app from this view.
- **Assets** — Lists all integrations and services registered on the app. Each asset shows its name, type (Action or AI), and permission level (Read or Write).

Publish an app

To publish an app:

1. When your app is ready for others to use, choose **Publish**.
2. Copy the shareable URL from the confirmation dialog.

You can continue editing after publishing. Changes do not affect the published version until you publish again.

Share an app

To share an app:

1. Choose **Share** to open the sharing controls.
2. Search for users by name, alias, or email.
3. Assign each person a role and choose **Share**.

Each shared user has a role that you can change at any time:

- **Co-owner** — Can edit, publish, and share the app.
- **Viewer** — Can view and interact with the published app.

Apps in Quick supports three access levels:

Access level	Description
Limited	Only users you explicitly share the app with can access it.
Account	Everyone in your Amazon Quick account can access the app.
Public	Anyone on the internet can access the app without signing in. Available on Free and Plus accounts only.

You can also toggle **Share with all** to make the app accessible to everyone in your Amazon Quick account, or choose **Copy link** to copy the published URL.

Publish an app to the public internet

To publish an app publicly (Free and Plus accounts):

1. Choose **Share** to open the sharing controls.
2. Set the access level to **Public**.
3. Choose **Publish**.

After publishing, the app receives a public URL that you can share with anyone. Anonymous viewers can use the app without creating an account.

Important

Public apps cannot use connectors, embedded visuals, embedded chat experiences, or Amazon Quick spaces. If your app uses any of these integrations, remove them before publishing publicly.

Manage public app usage

When your app is published publicly, AI inference usage counts against your subscription quota. Monitor your usage in the Amazon Quick account settings.

To manage or revoke public access:

1. Choose **Share** to open the sharing controls.

2. Change the access level from **Public** to **Limited** or **Account**.

Revoking public access takes effect immediately. The public URL stops working and anonymous viewers can no longer access the app.

What you can build with apps in Quick

You can build complete web applications with apps in Quick, not just static documents or simple forms. Your apps run in the browser, respond to user interaction, and connect to live data.

User interfaces

The agent generates responsive layouts using standard web components: navigation menus, tabs, cards, tables, charts, forms, modals, and buttons. You can describe the layout you want in natural language, or let the agent choose a layout based on the purpose you describe. Applications can include multiple views or sections, and the agent can implement navigation between them.

Role-based views are possible. If you tell the agent which users should see which content (for example, by alias or role), it can build conditional visibility into the interface using the user identity API available in the sandbox.

Data storage

Every app has access to built-in app storage, a key-value system that requires no external setup. There are two modes: private storage (scoped to the individual user, invisible to others) and shared storage (visible to everyone who accesses the app). Built-in storage scales to support large numbers of records.

You can also store and retrieve data through external systems using connectors or through Amazon Quick spaces. For more information, see [Working with data in apps in Quick](#).

AI inference

With apps in Quick, you can apply AI reasoning to features within your app. Common use cases include summarizing submitted text, classifying data into categories, generating recommendations, and answering natural language questions based on connected data.

AI inference supports two modes:

- **Free-form text** — Generate summaries, creative writing, chat responses, and any unstructured text output.
- **Structured output** — Extract structured data, classify inputs, and parse information into JSON. Uses tool-use patterns to guarantee valid, typed responses.

Available models range from fast (for simple tasks), to balanced (general purpose), to most capable (for complex reasoning). For more information about AI inference, see [Handling write approvals when AI inference is active](#).

Embedded visuals

With apps in Quick, you can embed live, interactive visuals from your Amazon Quick Sight dashboards directly in the app. Charts, tables, KPIs, and graphs update in real time with your data.

Each embedded visual must represent a unique dashboard, sheet, and visual combination. You can embed up to 8 visuals per app, with a minimum size of 100×100 pixels per visual. Embedded visuals are fully interactive and support filters, tooltips, and drill-downs. Viewers see visuals based on their Quick access permissions.

Embedded chat experiences

You can add a conversational chat interface to your app by connecting a Amazon Quick space as a knowledge base. The app creates its own agent instance that answers questions based on the documents and data in that space. This is useful for building knowledge portals, FAQ tools, or help desks within an app.

Important

You cannot currently embed a standalone chat agent that was created separately in Amazon Quick. The app's embedded agent is a new instance backed by the space you connect.

Public-facing applications

You can publish apps to the public internet so that anyone with the link can use them—no Amazon Quick account required. Public apps are available on Free and Plus accounts.

Public apps support the following capabilities:

- **Shared storage** — Anonymous viewers can read from and write to shared app storage. Private storage is not available to anonymous viewers.
- **AI inference** — You can enable AI inference for public apps. Usage counts against the app owner's subscription quota.
- **Custom domains** — Public apps are served from a unique URL on the Quick domain. Custom domains are not supported.

Public apps do not support connectors, embedded visuals, embedded chat experiences, or Amazon Quick spaces. For the full security model, see [Public app security](#).

File handling

With apps in Quick, you can generate files for download through the secure bridge API. Supported file types include CSV, JSON, plain text, HTML, XML, and Markdown. Apps can also read files from Amazon Quick spaces and display their contents. For more information about bridge API restrictions, see [Sandbox restrictions](#).

Important

Direct browser APIs for file creation are blocked by the sandbox. All downloads must go through the secure download utility provided by the platform.

Connecting to integrations in apps in Quick

You can connect your apps in Quick apps to external services and data sources through the integration framework in Amazon Quick. The same integration model used elsewhere in Quick applies in apps in Quick. The agent registers integrations on your app, and you and your users authorize them. For more information about integrations, see [Work with integrations in Amazon Quick](#).

Prerequisites

Before you connect integrations, confirm the following:

- You have an active Amazon Quick account with Author Pro or higher access.
- Your administrator has enabled the connectors you plan to use.

Integration types

Quick organizes integrations into two categories:

- **Knowledge integrations** — Connect to data sources that the system can search and reason over. In apps in Quick, this is most commonly a Amazon Quick space used as a knowledge base for an embedded chat experience.
- **Action integrations** — Connect to external applications and services where your app can read data, write data, or perform operations. Examples include Microsoft Outlook, SharePoint, OneDrive, Slack, Jira, and Asana. Quick provides over 100 pre-built action connectors. You can also create custom connectors using REST API, OpenAPI, or MCP specifications.

Authentication models

Action integrations support two authentication models:

- **User authentication** — The integration uses the credentials of the individual end user. Every user must have their own access to the connected resource.
- **Service authentication** — The integration uses a service-level credential configured by an administrator. All users of the app share that access. Not all connectors support service authentication.

Adding an integration to your app

You do not need to set up integrations manually. To add an integration:

1. Ask the agent to use a specific service in your instructions. For example: "Connect to my Outlook calendar and display my meetings for this week."
2. Authorize the integration when the agent prompts you.
3. Verify that the integration appears in the Assets view.

When other users open the published app, they authorize their own credentials for each registered integration on first use.

How apps in Quick works with other Quick capabilities

Apps in Quick connects to several other capabilities within Amazon Quick:

- **Amazon Quick spaces** — An app can read files from a space, write new files to it, and use a space as the knowledge base for an embedded chat experience. For more information, see [Organize, collaborate, and share resources with spaces in Amazon Quick](#).
- **Amazon Quick Flows** — There is no direct integration between apps in Quick and Quick Flows. Use a space as a shared data layer instead. A flow processes data and saves its output to a space. The app reads from that space and displays the results. For more information, see [Using Amazon Quick automations](#).

Integration tips

- **Outlook and OneDrive** — These connectors do not reliably resolve the me alias in API calls. Instruct the agent to use the user identity API to retrieve the current user's information instead.
- **Slack** — The Slack connector's API schema does not include all available Slack API operations. Provide specific details about the Slack API operations you need, such as the channel ID and message format.
- **Integration removal** — You cannot currently remove a registered integration from an app through the settings UI. As a workaround, ask the agent to remove references to the integration in the app code.

Tip

Limited integrations in public apps — Public apps do not support connectors, embedded visuals, embedded chat experiences, or Amazon Quick spaces. If you plan to publish your app publicly, design it to use only shared storage and AI inference.

Working with data in apps in Quick

You can persist and manage data within your apps in Quick apps in multiple ways.

Built-in app storage

The simplest way to persist data is with built-in app storage. It is a key-value system that requires no external setup and scales to support large numbers of records.

Storage scopes

- **Private storage** — Data visible only to the current user. Use for personal settings, preferences, saved filters, bookmarks, and per-user state.
- **Shared storage** — Data visible to anyone with app access. Use for collaborative lists, comments, votes, shared configuration, and team data.

Storage operations

Operation	Private	Shared	Description
Put Item	Yes	Yes	Store or update a key-value item.
Get Item	Yes	Yes	Retrieve an item by key.
List Items	Yes	Yes	List items with optional key prefix filter.
Delete Item	Yes	Yes	Remove an item by key.
List by Tag	No	Yes	Query shared items by tag (secondary index).

Key concepts

- **Tables** — Logical groupings for your data. You define table names; no setup is required.
- **Keys** — Each item has a unique key within its table. Keys can be up to 255 characters.
- **Values** — String values up to 350 KB. Use JSON serialization for structured data.
- **Tags (shared only)** — Optional categorization strings on shared items. Tags enable efficient querying by category.
- **Write modes** — UPSERT (default) overwrites existing items. INSERT fails if the key already exists, which is useful for preventing duplicates.

Note

Each app has completely separate storage. Data persists across user sessions and app reloads.

Note

In public apps, anonymous viewers can read from and write to shared storage but cannot access private storage. Design your data model accordingly if you plan to publish your app publicly.

Exporting data

You can ask the agent to add export functionality. Apps in Quick supports exporting data as CSV, JSON, PDF, and Excel files through the bridge API. You can also have your app write data snapshots to a connected space. This serves as a backup and makes the data available to other Quick capabilities.

Handling write approvals when AI inference is active

When an app uses AI inference and writes data (to shared storage or through an action connector), you must review and approve each write payload. This security measure ensures that you review AI-generated content before the app persists it.

Three strategies can reduce the frequency of approval prompts:

1. **Batch writes** — Collect all items and save them under a single storage key in one operation.
2. **Separate AI from writes** — Design the app so that AI processing and data persistence happen in distinct user actions.
3. **Remove AI inference when not needed** — If your app does not use AI-generated content, ensure the AI inference integration is not registered. Without AI inference, you can choose "Allow on this app" for write operations. The permission persists across sessions.

Accessing user information

With apps in Quick, you can access the current user's identity at runtime. This lets you personalize the app experience, display the user's name, or implement per-user logic.

Available user information includes email, first name, last name, and identity name.

Tip

Combine user identity with private storage to build personalized experiences such as greetings, saved preferences, bookmarks, and custom dashboards per user.

Security and sandbox in apps in Quick

Apps in Quick inherits the enterprise-grade authentication and authorization of Amazon Quick. Users access apps through their existing Quick identity, and every app runs inside a secure sandboxed iframe.

Authentication

- **Single sign-on** — Users authenticate through their organization's identity provider (SSO, Active Directory, IAM) via Quick. No additional credentials are needed.
- **Session security** — All app interactions run within an authenticated Quick session. Tokens are managed automatically by the platform.

Authorization layers

Layer	What it controls
App access	Who can view or edit the app (set by the app owner when sharing)
Integration approval	Which connectors, spaces, and dashboards the app can access (set by the author during authoring)
Runtime permissions	What data and actions are available to the viewer based on their own Quick permissions
Connector auth	How the connector authenticates with the external API (configured by admin)
Public access	Whether anonymous viewers can access the app without signing in (set by the app owner when sharing; Free and Plus accounts only)

Important

App viewers can only access data they are already authorized to see in Quick. Embedding a dashboard visual does not bypass row-level security or column-level permissions.

Integration consent model

When the apps in Quick agent adds an integration to your app (connector, space, dashboard visual, or AI inference), it prompts you for approval. This consent model ensures:

- You know exactly what external calls your app makes.
- You control READ vs WRITE permissions.
- The published app never includes unapproved integrations.
- App viewers inherit the approved integration scope, not broader access.

Sandbox restrictions

Every apps in Quick app runs inside a sandboxed iframe with strict security policies. The sandbox only permits script execution. All other capabilities (navigation, popups, direct network access) are restricted.

- **Link navigation** — Apps cannot open external URLs directly. Users can follow links by pressing Cmd+Click (macOS) or Ctrl+Click (Windows).
- **External resources** — The Content Security Policy blocks loading images, scripts, fonts, and other assets from external servers. Use inline SVG graphics, Base64-encoded image data, or image files loaded from a Amazon Quick space.
- **Network requests** — App code cannot make direct HTTP requests to external servers. All communication with external systems goes through the secure bridge API or a registered action connector.
- **File downloads** — File downloads must use the `downloadFile` function from the apps in Quick runtime library.
- **Public app isolation** — Public apps run in the same sandbox as private apps but cannot access connectors, embedded visuals, embedded chat experiences, or Amazon Quick spaces. Only shared storage and AI inference are available to anonymous viewers.

Public app security

Public apps allow anonymous access without authentication. The following security measures apply:

- **No identity** — Anonymous viewers have no user identity. The user identity API returns null values for public viewers.
- **No private storage** — Anonymous viewers cannot access private storage. Only shared storage is available.
- **No integrations** — Public apps cannot use connectors, embedded visuals, embedded chat experiences, or Amazon Quick spaces.
- **Rate limiting** — AI inference requests from public apps are rate-limited to prevent abuse. Usage counts against the app owner's subscription quota.
- **Same sandbox** — Public apps run in the same sandboxed iframe with the same Content Security Policy as private apps.

Tips for building effective apps

These tips come from patterns observed in successful app-building sessions.

- **Start with a detailed first prompt** — The first prompt sets the foundation. Including the purpose, audience, data model, and integration needs in your initial description gives the agent a clear starting point and reduces rework.
- **Make changes one at a time** — The agent produces more reliable results with focused, single-purpose instructions than with broad requests that modify many things at once.
- **Provide context about your goals** — When you explain why you want a change, not just what you want, the agent can make better implementation choices.
- **Investigate before changing** — The agent executes code changes with every instruction. If you want to explore options without modifying the app, explicitly say so: "Investigate the current data flow but do not write any code yet."
- **Note stable version numbers** — After completing a working feature, note the current version. If subsequent changes cause problems, you can restore the stable version immediately using the version selector.

- **Add a debug panel for complex apps** — For apps with significant logic, a debug panel that logs storage operations, integration calls, and state changes makes it easier to track down problems. You can make the panel visible only to yourself based on your user alias.
- **Specify roles when your app has different user types** — If different people should see different things, include that in your instructions.
- **Use structured AI output** — When you need data in a specific format (JSON, categories), use structured output mode instead of asking for free-form text. It guarantees valid, parseable responses.
- **Scope storage thoughtfully** — Use private storage for user-specific data and shared storage for collaborative features. Keep table names descriptive.
- **Approve integrations deliberately** — Review each integration approval carefully. Only approve connectors and spaces your app actually needs.
- **Design for anonymous viewers for public apps** — If you plan to publish your app publicly, remember that anonymous viewers have no user identity and cannot access private storage. Use shared storage for all data that public viewers need, and avoid features that depend on user identity.

Limitations of apps in Quick

The following limitations apply to apps in Amazon Quick in its current release.

Regional availability

Apps in Amazon Quick is available in a subset of the AWS Regions that Amazon Quick supports: US East (N. Virginia), US West (Oregon), Europe (Ireland), and Asia Pacific (Sydney). Support for additional Regions is being added.

Limits

Limits

Resource	Limit
Dashboard visuals per app	8
Visual minimum size	100 × 100 pixels

Resource	Limit
Storage value size	350 KB per item
Storage key length	255 characters
Table name length	255 characters

Building experience

- **Content policy violations** — Prompts occasionally trigger security guardrails as false positives. When this happens, the agent cannot process further prompts in the current session. Close the editor and wait 15 to 20 minutes for the session to time out, then reopen the app.
- **Agent memory in long conversations** — The agent's accuracy decreases as the conversation grows longer. For complex applications, consider building in phases.
- **No image upload from agent chat** — You cannot upload screenshots or reference images to the agent.
- **No investigation-only mode** — The agent executes code changes with every instruction. To have the agent investigate without making changes, explicitly say "Investigate but do not write code yet."
- **Network sensitivity** — The editor uses a WebSocket for real-time streaming. Unstable connections can cause prompts to fail silently.

Sharing and access

- **Subscription requirement** — Only users with Author, Professional, Author Pro, Enterprise, or Admin Pro subscriptions can view apps.
- **External access** — Users must have a Amazon Quick account to view an app, unless the app is published publicly. Public access is available on Free and Plus accounts only. Public apps cannot use connectors, embedded visuals, embedded chat, or spaces.

Portability

- **No code export** — You cannot download the source code of an app.

- **Duplication** — You can duplicate an app to create a copy.

Integration limitations

- **No direct Quick Flows integration** — Use a space as an intermediary.
- **No standalone agent embedding** — You can create a new agent instance backed by a space.
- **No integration removal** — You cannot remove a registered integration from an app through the settings UI.

Sandbox restrictions

For detailed information about sandbox restrictions, see [Sandbox restrictions](#). Key restrictions include:

- **No direct link navigation** — Users must use Cmd+Click or Ctrl+Click.
- **No external images** — CSP blocks loading images from external URLs.
- **No built-in app analytics** — As a workaround, you can ask the agent to implement a view counter using shared app storage.

Troubleshooting apps in Quick

This section describes common issues and how to resolve them.

Recovering from content policy violations

If you receive the error "This request could not be processed due to content policy violations":

1. Close the apps in Quick editor.
2. Wait 15 to 20 minutes for the session to reset.
3. Reopen the app.
4. Try a simpler, more focused prompt.

All prompts in the same session after a violation continue to fail until the session resets.

Recovering from a broken app

If an app stops loading or displays errors after an edit:

1. Use the version selector to restore a previous working version.
2. If the editor is not accessible, close the browser tab, wait 15 to 20 minutes, and reopen.
3. In a fresh session, try: "Review all app code. Check the existing code against the actual types and interfaces exported by the runtime library. Fix any issues."
4. If an integration change caused the issue, try: "The prior issue impacting integrations has been resolved. Please try again."

Fixing connector authentication issues

If you see a "Sign in to connectors" dialog with no sign-in option, one or more integrations were not properly registered during app creation. In the editor, ask the agent to retry registering each affected connector individually. After successful re-registration, publish the updated app.

Debugging complex apps

You can ask the agent to add a debug panel that logs storage operations, integration calls, and state changes. Make the panel visible only to yourself (based on your user alias) to keep it hidden from end users.

File export not working

If an export feature does not work, ask the agent: "Use the `downloadFile` function from the runtime library to handle the file export." The sandbox blocks direct browser APIs for file creation. For more information, see [Sandbox restrictions](#).

Work with integrations in Amazon Quick

Integrations in Amazon Quick connect you to external applications and services for AI-powered analysis and automation. With integrations, you can perform actions in external applications and bring data from external applications to enhance your AI agents and chat experiences.

You can create three types of integrations:

Connectors

Initiate real-time calls to perform actions in external services. Use action connectors to address user natural language chat requests, automate workflows, send notifications, and/or trigger processes in connected applications.

Data access integrations

Connect Quick with external applications to access data. Data access integrations serve as the foundation for creating knowledge bases. You cannot use these integrations directly for analysis.

Knowledge bases

Connect to data access integrations to make specific content accessible for AI analysis. Choose which data or knowledge to include from external applications. Add knowledge bases to spaces or use them directly with chat agents.

Use the Integrations tab in the Amazon Quick console to create these connections. You can connect to Model Context Protocol (MCP) servers, Google Drive, Amazon S3, web crawlers, OneDrive, and other supported applications.

Note

Not all applications support all integration types. Available options depend on each application's capabilities.

Topics

- [Supported integrations](#)
- [Set up integrations in the console](#)

- [Integration workflows](#)
- [Knowledge bases](#)
- [Connectors](#)
- [Integration-specific guides](#)
- [Bring Your Own Amazon Q Business Index \(BYOI\)](#)
- [Bring your own Amazon Bedrock managed knowledge base](#)

Supported integrations

Amazon Quick supports integrations with various third-party applications and services. Each integration supports different combinations of actions and knowledge base creation capabilities. The following table shows the supported integrations and their capabilities.

Supported Integration Capabilities

Integration	Actions	Knowledge Base
Adobe Marketing Agent	✓	—
Airtable	✓	—
Amazon S3	✓	✓
Asana	✓	—
Atlassian Confluence Cloud	✓	✓
Atlassian Jira Cloud	✓	—
AWS Agent Registry	✓	—
BambooHR	✓	—
Box	✓	—
Canva	✓	—
Cisco Webex Meetings	✓	—

Integration	Actions	Knowledge Base
Cisco Webex Video Messaging	✓	—
Dropbox	✓	—
Dun & Bradstreet	✓	—
GitHub	✓	—
Gmail	✓	—
Google Analytics	✓	—
Google Calendar	✓	—
Google Chat	✓	—
Google Docs	✓	—
Google Drive	✓	✓
Google Meet	✓	—
Google Sheets	✓	—
Google Slides	✓	—
HG Insights	✓	—
HubSpot	✓	—
HuggingFace	✓	—
Intercom	✓	—
Linear	✓	—
Microsoft OneDrive	✓	✓
Microsoft OneNote	✓	—

Integration	Actions	Knowledge Base
Microsoft Outlook	✓	—
Microsoft SharePoint Online	✓	✓
Microsoft Teams	✓	—
Model Context Protocol (MCP)	✓	—
Monday.com	✓	—
Moodys GenAI Ready Data	✓	—
New Relic	✓	—
Notion	✓	—
OpenAPI Specification	✓	—
PagerDuty	✓	—
QuickBooks	✓	—
REST API	✓	—
Salesforce	✓	—
SAP Bill of Materials	✓	—
SAP Business Partner	✓	—
SAP Material Stock	✓	—
SAP Physical Inventory Docs	✓	—
SAP Product Master	✓	—
ServiceNow	✓	—
Shopify	✓	—

Integration	Actions	Knowledge Base
Slack	✓	—
Smartsheet	✓	—
Snowflake Cortex Agent	✓	—
Visier Agent	✓	—
Web Crawler	—	✓
WhatsApp	✓	—
Zapier	✓	—
Zendesk Suite	✓	—
Zoom	✓	—
ZoomInfo	✓	—

Note

Not all applications support all integration types. The available options depend on the capabilities of each specific application and your user role.

Integration capability definitions

Actions

Call APIs and perform actions in external applications directly from Amazon Quick. You can share connectors with other users and use them in automated workflows.

Knowledge base

Create searchable repositories of information from external sources. Knowledge bases are children of data access integrations. Add them to spaces or use them directly in chat agents.

Authentication method definitions

User auth

Custom user-based OAuth authentication requiring base URL, client ID, client secret, token URL, auth URL, and redirect URL.

Service auth

Service-to-service authentication using either API key (with base URL and email) or service-to-service OAuth (with base URL, client ID, client secret, and token URL).

Managed OAuth 2.0

Managed OAuth 2.0 authentication flow with provider-specific sign-in interface.

AWS credentials

AWS-specific authentication using AWS access keys and permissions.

Basic auth

Username and password authentication.

Form/SAML auth

Form-based or SAML authentication with configurable field selectors.

JSON schema

Schema-based authentication for OpenAPI specifications.

Set up integrations in the console

The console organizes integrations into separate categories based on their purpose. Use **Knowledge** to connect data sources for Q&A and insights. Use **Connectors** to set up connectors that perform operations in external applications. The setup process adapts based on the integration you select, your subscription, and existing integrations.

Choose integration options

When you set up an integration, the console guides you based on several factors:

- **Integration capabilities** – Each application supports different combinations of actions and knowledge base creation. For example, Google Drive supports both actions and knowledge base creation. Web Crawler supports knowledge base creation only.
- **Subscription** – Configuring integrations requires an Enterprise subscription. This includes creating connectors, setting up knowledge bases, and managing integration settings. Users with a Professional subscription can use integrations that have been shared with them.
- **Existing integrations** – When you choose a connector that already exists, the console shows your existing connectors before offering to create new ones.

View setup process examples

The following examples show how different integrations guide you through different console setup processes.

Google Drive – Set up a knowledge base

Google Drive supports knowledge base creation through user-managed or admin-managed authentication.

1. In the console, choose **Knowledge**.
2. Find **Google Drive** and choose the **Add (+)** icon.
3. Choose your authentication method and complete the sign-in flow.
4. Enter a name and description for your knowledge base.
5. Select the files and folders you want to index, then choose **Create**.

Google Drive – Set up an action connector

Google Drive also supports connectors for performing file operations directly from .

1. In the console, choose **Connectors**.
2. Choose the **Create for your team** tab.
3. Find and choose **Google Drive**.
4. Enter a **Name** for your connector. Optionally, choose **+ Add Description** to add a description.
5. Choose your **Connection type** and **OAuth Configuration**, then complete the authentication setup.

6. Review the available actions and choose **Publish**.

Available connectors tab

The **Available** tab on the Connectors page shows connectors that use managed OAuth and are ready to use without additional configuration. Users can connect directly by choosing **Connect** on the connector card.

1. In the console, choose **Connectors**.
2. On the **Available** tab, find the connector you want to use.
3. Choose **Connect**.
4. Complete the sign-in flow for the application.

Working with existing connectors

When you choose a connector on the **Create for your team** tab that already exists, the console displays your existing connectors before offering to create a new one.

1. In the console, choose **Connectors**.
2. Choose the **Create for your team** tab.
3. Choose a connector that you've configured before.
4. A dialog appears showing your existing connectors. To use an existing connector, choose it. To create a new one, choose **No, create new**.

Integration management options

After creating integrations, you can manage them through several console options:

- **Edit integration** – Modify integration settings, authentication details, and configuration options.
- **Delete integration** – Remove integrations with confirmation dialogs to prevent accidental deletion.
- **Knowledge base management** – Separate flows for creating, editing, and deleting knowledge bases associated with your integrations.

Integration workflows

The following procedures describe the general workflows for creating and managing different types of integrations in Amazon Quick.

Creating a knowledge base

Knowledge bases connect external data sources to Amazon Quick so you can ask questions about and get insights from your content.

1. In the Amazon Quick console, choose **Knowledge**.
2. Find the application you want to connect to (for example, **Google Drive**, **Microsoft OneDrive**, or **Amazon S3**) and choose the **Add (+)** icon.
3. Complete the authentication process for your chosen application.
4. Enter a **Name** and optional description for your knowledge base.
5. Select the files and folders you want to include in your knowledge base using the file picker or sync options.
6. Choose **Create**.

After you choose **Create**, the data sync starts automatically.

Creating a connector

Connectors enable you to perform actions in external applications directly from Amazon Quick.

1. In the Amazon Quick console, choose **Connectors**.
2. Choose the **Create for your team** tab.
3. Find and choose the application you want to connect to (for example, **Google Drive**, **Atlassian Confluence Cloud**, or **Slack**).
4. Enter a **Name** for your connector. Optionally, add a description.
5. Configure the connection type and authentication method. For more information about authentication options, see [the section called "Authentication methods"](#).
6. Review the available actions and choose **Publish**.

After successful creation, your connector is available for use in Amazon Quick chat, workflows, and Amazon Q Apps.

Managing existing integrations

You can edit, delete, share, and manage existing integrations from the Amazon Quick console.

To edit an integration

For knowledge bases:

1. In the Amazon Quick console, choose **Knowledge**.
2. Choose the knowledge base you want to edit.
3. Choose the menu icon (:) and select **Edit**.
4. Modify the settings as needed and choose **Save changes**.

For connectors:

1. In the Amazon Quick console, choose **Connectors**.
2. Choose the connector you want to edit.
3. Choose the menu icon (:) and select **Edit**.
4. Modify the connector settings as needed and choose **Save changes**.

To delete an integration

For knowledge bases:

1. In the Amazon Quick console, choose **Knowledge**.
2. Choose the knowledge base you want to delete.
3. Choose the menu icon (:) and select **Delete**.
4. In the confirmation dialog, choose **Delete** to confirm.

For connectors:

1. In the Amazon Quick console, choose **Connectors**.
2. Choose the connector you want to delete.
3. Choose the menu icon (:) and select **Delete**.
4. In the confirmation dialog, choose **Delete** to confirm.

The integration is permanently removed from your account. Any dependent resources (such as knowledge bases) that rely on this integration will be impacted.

To share an integration

From the connector or knowledge base details page, choose the menu icon (⋮) and select **Share**.

Managing connectors

Connectors have additional management options available from the connector details page.

Sign in or re-connect

For connectors that use Default OAuth app or Custom OAuth app authentication, you must sign in to the service before you can use its actions. If you have not yet signed in, a **Sign in** button appears on the details page. After you sign in, you can use **Re-Connect** to re-authenticate if your session expires or the connection is interrupted.

Test action APIs

Choose **Test action APIs** in the Actions section of the details page to test individual actions provided by the connector. This lets you verify that the connection is working correctly and that the service responds as expected.

Note

The **Test action APIs** option is available for connectors only. Knowledge bases do not support action testing.

Knowledge bases

A knowledge base is an organized, indexed collection of documents or content from data sources optimized for generative AI-powered retrieval and question answering. Whether your team stores documentation in Confluence, collaborates through SharePoint, or manages files in cloud storage, you can bring all this information into one unified search experience by creating knowledge bases.

The built-in integrations can be set up with just a few clicks to sync your data in Quick and make it effortless to tap into your organization's knowledge across Google Drive, OneDrive,

Confluence, SharePoint, S3, and Web Crawler. Whether your team stores documentation in Confluence, collaborates through SharePoint, or manages files in cloud storage, you can bring all this information into one unified search experience with the help of knowledge bases.

How knowledge bases work

Knowledge base is an indexed collection of documents or content from data sources such as Google Drive, optimized for generative AI-powered retrieval and question answering. Multiple knowledge bases can be created from the same source, and all can reside within a shared Quick Index. For example, if you sync two folders from Google Drive and create two knowledge bases — one for “Policy Documents” to answer queries such as *“What’s our refund policy”* and one for “Customer feedback” to answer queries such as *“What are the common customer complaints”* — both can be part of the same index. Quick distinguishes between them using the knowledge base id, so queries can be filtered to retrieve only the relevant documents from the desired knowledge base. This allows users to organize, secure, and retrieve information relevant to different domains or use cases, even though the underlying data is indexed together.

Your knowledge bases can be used individually or shared with team members through Amazon Quick spaces. Our coarse-grained access control enables security at the knowledge base level, ensuring that users only receive information from knowledge bases they're authorized to access.

Creation process

You can create knowledge bases while setting up a new data access integration and use existing integrations to create additional knowledge bases:

1. **Data access integration setup** - Connect to your external data source
2. **Content selection** - Choose which content to include through filters and scope settings
3. **Indexing** - Amazon Quick processes and indexes the selected content
4. **Availability** - The knowledge base becomes available for use in spaces and by AI agents

Capabilities

Each knowledge base provides the following capabilities:

- **Content indexing** - Processes text, documents, and structured data from external sources
- **Semantic search** - Enables AI-powered search across indexed content

- **Automatic synchronization** - Keeps content up-to-date with configurable sync schedules
- **Coarse-grained access control** - Ensures that users only receive information from knowledge bases they're authorized to access.
- **Multi-space usage** - Can be used across multiple spaces and by different AI agents

General workflow

The typical workflow for working with knowledge bases follows these steps:

1. **Set up data source integration** - Connect to your external application (such as SharePoint, Google Drive, or Confluence) with appropriate authentication. For more information, see [the section called "Integration-specific guides"](#).
2. **Create a knowledge base** - You can create a knowledge base while configuring your new integration. Configure your content filters by setting up include filters, file type restrictions, and folder selections to focus on relevant content.
3. **Set sync schedule** - Data refresh frequency is set to daily by default. You can edit the sync frequency to configure how often the knowledge base should be updated with new content from the source.
4. **Monitor and manage** - Review sync status, manage access permissions.

Common configuration settings

Knowledge bases share common configuration patterns across different data source integrations. Understanding these settings helps you optimize content indexing and manage sync behavior effectively.

Note

While these configuration options are available across most integrations, specific settings and available options may vary depending on your chosen data source integration.

Service principal for knowledge base operations

If you use a customer-managed key (CMK) as your default CMK and Q data key (see [the section called "Managing AWS KMS keys"](#)), Amazon Quick accesses your AWS KMS key using the

`qbusiness.amazonaws.com` and `quicksight.amazonaws.com` service principals during knowledge base sync operations.

Note

Ensure that any policies governing access to your CMK permit both the `qbusiness.amazonaws.com` and `quicksight.amazonaws.com` service principals.

File size and content limits

Configure file size limits to optimize processing performance and manage storage costs. The specific limits vary by content type and are displayed in the console when you configure your knowledge base.

Standard text documents

Applies to documents like PDFs, Word files, and text files. File size limit is 500 MB.

Video files

Available when video processing is enabled. Supported formats include `.mp4`, `.mov`, `.m4v`. File size limit is 10 GB (10240 MB). Quick Index supports up to **10 video files per GB of storage**. If your use case requires higher video volumes, please open a ticket with AWS support to extend this limit.

Audio files

Available when audio processing is enabled. Supported formats include `.mp3`, `.wav`, `.m4a`, `.flac`, and `.ogg`. Limit is 2 GB (2048 MB) for audio files.

Note

Files with extracted text that exceeds the 30 MB system limit are not indexed, regardless of the original file size. The maximum amount of text that can be extracted from a single document is 30 MB.

Images

Quick Index applies the following limits for images:

- **Per-document limit:** 500 images per document
- **Per-GB limit:** 10K images per GB of index storage
- **Per-index limit:** 2M images per index

If your use case requires higher image volumes, please open a ticket with AWS support to extend these limits.

Sync schedule and safeguards

Configure how often your knowledge base updates and protect against unintended content deletion:

Sync frequency

Data refresh frequency is set to daily by default. You can edit the sync frequency to configure how often the knowledge base should update with new content from the source

Document deletion safeguard

Protect your indexed content from accidental mass deletion by setting a maximum deletion percentage threshold. If a sync job would delete more documents than your threshold allows, the deletion phase is skipped, preserving your existing indexed content.

This safeguard protects against temporary network issues, permission changes, or source system problems that might make content temporarily unavailable.

Maximum sync duration

Each sync run has a maximum runtime of 14 days. If a sync run is still in progress after 14 days, Amazon Quick ends the run with a status of **FAILED** and the following error message:

Maximum sync time exceeded: Please edit your sync configuration to reduce the number of files to be crawled and ingested. For web crawler data sources, we recommend using Web Search feature if you want to chat with public websites like wikipedia.org.

If your sync run reaches this limit, edit your knowledge base to narrow the sync scope. For example, apply include or exclude filters, restrict the file types or folders being crawled, or split the content

across multiple knowledge bases. For web crawler data sources, consider using the Web Search feature instead when your goal is to chat with large public websites.

Best practices for managing ACLs in knowledge bases

With document-level access control lists (ACLs), Amazon Quick enforces source-document permissions for an ACL-aware knowledge base. Each authorized user retrieves only the indexed documents they have permission to access. Use ACLs when different users need access to different documents in the same knowledge base.

You're responsible for keeping source identities, groups, and document permissions accurate. Quick enforces the synchronized document permissions on every retrieval. For supported integrations, it also verifies document access against the source in real time when returning results. If Quick can't evaluate document permissions for a query, it returns no documents rather than unfiltered results.

Quick synchronizes identity and document-permission changes on the knowledge base refresh schedule, which is every 24 hours by default. Configure a different schedule when your access-change requirements demand it.

Sharing and document access are separate controls

Sharing a knowledge base and granting document access are separate controls. Knowledge-base sharing determines who can use the knowledge base. For ACL-aware knowledge bases, source-document ACLs further limit which indexed documents each authorized user can retrieve. Review both controls before granting access.

For more information about configuring ACLs for a specific data source, see [Amazon S3](#), [Google Drive](#), or [Microsoft SharePoint](#). For [Atlassian Confluence Cloud](#) and [Microsoft OneDrive](#), configure document-level ACLs in the console where available.

To verify document-level access controls and troubleshoot permission issues, see [the section called "Check document access \(ACL verification\)"](#).

Note

Quick treats all email addresses as case-insensitive. `JohnDoe@example.com`, `johndoe@example.com`, and `JOHNDOE@example.com` are all considered the same user.

Plan ACL-aware knowledge bases before creation

Before you create an ACL-aware knowledge base, complete the following steps:

1. Confirm that your integration supports document-level ACLs.
2. Confirm the identity attribute that Quick uses to resolve users and groups.

Quick resolves ACLs within the namespace of the knowledge base creator. For details, see [Limitations](#).

3. Remove shared or recycled identities from source ACLs before you assign them to another person.
4. Select a refresh schedule that meets your access-change requirements. For Amazon S3, permission changes take effect at the next sync, so plan the schedule accordingly.
5. Test document access with representative users before you share the knowledge base broadly. To check document access, see [the section called "Check document access \(ACL verification\)"](#).
6. Confirm that the knowledge base isn't required for Quick Research.
7. Assign at least one additional owner to an admin-managed knowledge base so that it remains manageable when its original creator leaves.

Important user management scenarios

Understanding email binding

Email addresses are bound to Quick users dynamically when users initiate chat interactions. This binding follows a first-come-first-serve approach. The first user to chat with a given email address establishes the binding for that identity within the namespace.

When an employee leaves your organization

When an employee leaves, clean up their access promptly:

1. Update the ACL configuration files to remove references to their email address. For example, in Amazon S3, update the global ACL file or metadata files.
2. Refresh the knowledge bases to apply the changes.

This prevents potential security issues if the email is later reassigned to someone else.

Updating knowledge base ACLs is separate from removing the user from Quick. For the full model of how user removal affects a user's assets and data, see [User lifecycle and data handling in Amazon Quick](#).

Share admin-managed knowledge bases with co-owners

Admin-managed knowledge bases (service credentials) are often used across teams and organizations. If the original creator leaves the company and no co-owners exist, the knowledge base becomes unmanageable — no one can edit settings, trigger syncs, or update permissions. To avoid this, share admin-managed knowledge bases with at least one additional owner. For more information, see [the section called "Sharing knowledge bases and data sources"](#).

When an email address is reassigned to a new employee

- ACL-aware knowledge base access is automatically locked for the reassigned email address to protect data security.
- Contact Quick support to clean up the previous user's access before the new employee can access documents associated with that email.

Limitations

When configuring document-level ACLs for your knowledge bases, be aware of these limitations:

- **Document-level ACL configuration is permanent** – You cannot turn on ACLs for a knowledge base created without ACL support. You also cannot turn them off after you turn them on. To change ACL configuration, create a new knowledge base with your desired setting from the start.
- **Shared email addresses within a namespace** – If multiple Quick users share the same email address within a namespace, the system denies access to everyone using that shared email. This safeguard prevents accidentally granting document access to the wrong person.
- **ACL resolution scope** – Quick resolves all ACLs within the namespace of the knowledge base creator. This applies whether you specify ACLs by email address or group name. Quick looks up identities in the creator's organizational context to ensure consistent identity resolution.
- **Email address recycling timing** – If your organization reassigns an email address from one employee to another, there's an important timing consideration. If the previous employee never used Quick for chat or AI interactions, and the email is reassigned before the next ACL refresh, the new employee may temporarily access documents intended for the previous employee.

To avoid this, complete the following steps in order:

1. Update your ACLs (if applicable, such as in Amazon S3) to remove the old user and add the new user.
2. Manually refresh your knowledge base, or wait for the automatic daily refresh.
3. Assign the email address to the new employee.

This ensures access permissions are properly synchronized before the new user begins using Quick.

Research compatibility

Knowledge bases with document-level ACLs enabled aren't currently compatible with Quick Research. If you need to use documents from an ACL-enabled knowledge base for research, create a separate knowledge base without ACLs for those documents.

Sharing knowledge bases and data sources

Amazon Quick enables you to share knowledge bases and data source connections with other users in your account. You can collaborate with multiple users on knowledge bases or reuse existing data source connections without re-entering credentials.

Note

All knowledge bases can be shared with the Viewer role, which allows users to query the knowledge base. The Owner role (co-owner sharing) is available only for admin-managed Microsoft SharePoint Online and Google Drive knowledge bases. Data source sharing is also limited to admin-managed SharePoint and Google Drive connections.

Sharing knowledge bases

Knowledge base owners can share their knowledge bases with other users and assign one of two roles:

Knowledge base sharing roles

Role	Permissions
Owner	View all tabs (Summary, Sync History, Settings, Permissions). Edit knowledge base details. Trigger manual syncs. Update the sync schedule. Share the knowledge base with others. Delete the knowledge base. Owners cannot edit the underlying data source connection unless they are also an Owner of the data source.
Viewer	View the knowledge base in the list and access the Summary tab only. Viewers cannot edit, share, delete, or trigger syncs.

Before you share, note the following:

- All knowledge bases support sharing with the Viewer role. Viewers can query the knowledge base.
- Any owner of the knowledge base can share it, not only the original creator.
- Co-owner sharing (the Owner role) is available only for admin-managed SharePoint and Google Drive knowledge bases. All other knowledge base types can be shared only with the Viewer role.
- The original creator of a knowledge base always retains the Owner role and cannot be removed.

Sharing a knowledge base

You can access the sharing permissions from either of the following locations:

- From the knowledge base list, choose the actions menu (:) next to the knowledge base and choose **Share knowledge base**.
- From within a knowledge base, choose the **Permissions** tab.

To add a user:

1. Choose **Add users & groups**.
2. In the search field, enter the user's name or email. Choose the user from the results, and then choose **Add**.

3. The user appears in the permissions list with a default role of **Viewer**. To change the role, open the **Permissions** dropdown next to the user's name and choose **Owner** or **Viewer**.

Changing or revoking knowledge base access

To change or remove a user's access to a knowledge base, complete the following steps.

1. Open the knowledge base and choose the **Permissions** tab.
2. Locate the user in the permissions list.
3. To change their role, open the **Permissions** dropdown and choose **Owner** or **Viewer**.
4. To remove access, choose **Revoke access** in the **Actions** column.

Sharing data source connections

Amazon Quick administrators can share data source connections with other users in their account. This enables multiple users to create knowledge bases using the same connection without re-entering credentials.

Before you share, note the following:

- Only users with Amazon Quick administrator privileges can share data source connections.
- Data source sharing is available only for admin-managed (service credentials) connections. User-managed connections that are created through single sign-on cannot be shared.
- If you revoke a user's access to a data source, the user's existing knowledge bases continue to sync. Knowledge bases that were previously created with the shared data source continue to sync normally.
- If an administrator deletes a data source, then existing knowledge bases that are linked to it remain visible. However, syncs fail with an error that indicates the underlying data source was deleted.
- Unlike knowledge bases, the original creator of a data source can be removed or have their role changed.

Data source sharing roles

Role	Permissions
Owner	Use the shared data source to create knowledge bases and edit the connection details (such as App ID or App Secret).
Viewer	Use the shared data source to create knowledge bases but cannot edit the connection details.

Sharing a data source connection

To share a data source connection with other users, complete the following steps.

1. In Amazon Quick, choose **Manage account** from the left navigation pane.
2. Choose **Manage assets**.
3. Under **Browse assets**, choose **Data sources**.
4. Select the checkbox next to the data source that you want to share.
5. Choose **Share**.
6. In the **Share assets** dialog, choose a **Permissions** level (**Owner** or **Viewer**).
7. In the **User or Group** field, enter the user's name or email and choose the user from the results.
8. Choose **Share**.

The user now sees the shared data source in the **Connected account** dropdown when the user creates a new knowledge base.

Editing a data source connection

Data source Owners can edit the connection details from the knowledge base list:

1. In the knowledge base list, choose the actions menu (:) next to a knowledge base that uses the data source.
2. Choose **Edit integration**.
3. Update the connection details and choose **Save**.

Important

If you edit a shared data source connection, then all knowledge bases that use it are updated. Review the impact before you save changes.

Changing or revoking data source access

To view or manage permissions for a data source, complete the following steps.

1. Choose **Manage account**, then choose **Manage assets**, and then choose **Data sources**.
2. Choose the data source name to open its details page.
3. Under **User & Groups Permissions**, locate the user and change their role or remove their access.

Sync reports and observability

With Amazon Quick, you can view sync schedules, sync activity logs, and detailed sync reports for every knowledge base. If your knowledge base has document-level access control (ACL) enabled, you can also verify individual document permissions.

Sync schedules and activity

To view sync schedules and activity, open a knowledge base and choose the **Sync schedules** tab.

Sync schedule

Displays the configured sync frequency and timing. The table shows the occurrence (for example, Daily), the start time for the next scheduled sync, and the email notification status. You can add or modify sync schedules. To receive email notifications for sync failures, turn on **Email on sync failure**.

Sync activity

Displays a historical log of all past syncs. You can filter by time frame (for example, Last 90 days) and sync status (for example, All, Completed, Failed).

The following table describes the sync activity columns:

Sync activity columns

Column	Description
Start time	The date and time the sync began.
Sync status	The outcome of the sync: Queued (waiting to run), Completed (all items processed), Completed with issues (some items were skipped or failed), or Failed.
Duration	How long the sync took to complete.
Total items scanned	The number of items processed during the sync.
Type	Whether the sync was Scheduled (automatic) or Manual (initiated by choosing the Sync now button).
Sync report	A View report link that opens the detailed sync report for that run.

Tip

To start a sync on demand, choose **Sync now** on the knowledge base detail page.

Review sync reports

To view detailed results for a specific sync, open a knowledge base and choose the **Sync reports** tab.

Use the **Select sync** dropdown to choose the sync run you want to inspect.

Overview

The overview section provides a summary of the selected sync. The following table describes the overview fields:

Sync report overview fields

Field	Description
Sync status	The overall outcome (for example, Completed, Completed with issues).
Sync type	Whether the sync was Scheduled or Manual.
Sync duration	Total time elapsed for the sync.
Sync start time	When the sync began.
Sync end time	When the sync finished.

Sync items

The sync items section shows all items that were processed during the sync, organized into the following categories:

Available items

Items that are present and indexed in the knowledge base after the sync:

- **Unmodified** – Items already indexed that have not changed since the last sync.
- **Modified** – Previously indexed items that were updated and re-indexed.
- **Added** – New items indexed for the first time.

Unavailable items

Items that are not available in the knowledge base after the sync:

- **Skipped** – Items filtered out (for example, empty files, unsupported formats, or items outside the date filter range).
- **Failed** – Items that could not be processed because of an error.
- **Deleted** – Items previously indexed that have been removed from the source.

Filter and inspect individual items

The **All items** tab displays the complete list of items that were processed during the sync. You can narrow the results by using the following filters:

Sync item filters

Filter	Description
Filter by status	Filter items by their sync status (for example, Added, Modified, Unmodified, Skipped, Failed, Deleted).
Filter by error type	Filter failed or skipped items by their specific error type.
Filter by item title	Search for a specific item by name.

After you set your filters, choose **Apply filters** to update the results.

The following table describes the sync item columns:

Sync item columns

Column	Description
Item Title	The name of the file, page, or item.
Item Status	The sync status for this item.
Error type	The category of error, if applicable.
Error message	A detailed description of why the item failed or was skipped.

Download report

To export the full sync results for offline analysis, choose **Download detailed report (.csv)** in the Sync items section. The CSV file includes all items with their status, error type, and error message.

Check document access (ACL verification)

If your knowledge base has document-level access control (ACL) enabled, you can verify whether a specific user has access to a particular document. Use this feature to troubleshoot when users cannot see documents that you expect them to access, or to confirm that sensitive documents are properly restricted.

 Note

ACL verification is available for admin-managed SharePoint and Google Drive knowledge bases with ACL management enabled. Amazon S3 knowledge bases with document-level ACLs are not included in this feature. Because the ACL configuration files are customer-managed, you can verify them directly in Amazon S3.

1. Open the knowledge base and choose the **Sync reports** tab.
2. In the sync items list, choose the actions menu (:) for an item.
3. Choose **View Access Details**.

The **Access Details** panel shows:

- The document name, status, and last synced timestamp.
- **Permission Checker** – Enter a user's email address and choose **Check Access** to verify whether that user can access the document.
- **Users and group membership** – Lists all users and groups with access to the document.

The following table describes the permission checker results:

Permission checker results

Result	Meaning
User has access	The user's identity matches the document's ACL. They see this document in Quick responses.
[user] does not have access to this document	The user is not in the document's ACL. Quick does not surface this document for them.
No access control list found	The ACL was not crawled for this document. Either ACL management is not enabled, or the required permissions are not configured correctly.

If you see "No access control list found", take the following steps:

- Verify that ACL management is enabled in the knowledge base settings.
- Confirm the app registration has the required ACL permissions for your data source.
- Run a full sync after you fix the permissions.

Monitor index storage usage

Index usage logs provide visibility into how your Amazon Quick index storage is consumed across knowledge bases and Spaces, including file uploads. You can use these logs to track growth trends, identify the largest sources, detect unexpected spikes, and plan capacity.

Before you query index usage data, you must configure INDEX_USAGE_LOGS delivery. For instructions on setting up delivery, see [the section called “Monitoring using CloudWatch Logs”](#).

Log schema

Each index usage log event includes the common fields that are shared across all Amazon Quick log types, such as `resource_arn`, `event_timestamp`, `log_type`, `account_id`, and `user_arn`. The event also includes the following index-specific fields:

Index usage log fields

Field	Type	Description
<code>consumed_index_size</code>	Integer	Total size (in bytes) consumed by the entire index. This is the authoritative total.
<code>source_type</code>	String	SPACE or KB.
<code>source_name</code>	String	Display name of the Space or knowledge base.
<code>source_arn</code>	String	Full ARN of the source.
<code>consumed_source_size</code>	Integer	Size (in bytes) consumed by this individual source.
<code>consumed_source_documents_count</code>	Integer	Number of documents in this source.

Note

Events are published per source on change. Not all sources emit events every day. To reconstruct the current state, use the most recent event per `source_arn`.

CloudWatch Logs Insights queries

The following CloudWatch Logs Insights queries help you analyze index usage data. All queries use the pattern `stats latest(field) by source_arn` to get the most recent state per source. The queries then aggregate as needed. Replace *YOUR_LOG_GROUP* with your log group name.

Size by source type (knowledge bases compared to Spaces)

Shows the total size split between knowledge bases and Spaces.

```
fields @timestamp, source_type, source_arn, consumed_source_size
| stats latest(consumed_source_size) as latest_size, latest(source_type) as type by
  source_arn
| stats sum(latest_size) as total_size by type
```

Total index size over time

Shows the index size trend by using the `consumed_index_size` field.

```
fields @timestamp, consumed_index_size
| stats latest(consumed_index_size) as total_index_size by bin(1d) as day
| sort day asc
```

Top knowledge bases by size

Shows the top 20 knowledge bases ranked by size.

```
fields @timestamp, source_type, source_arn, source_name, consumed_source_size
| filter source_type = "KB"
| stats latest(consumed_source_size) as latest_size, latest(source_name) as name by
  source_arn
| sort latest_size desc
| limit 20
```

Top Spaces by size

Shows the top 20 Spaces ranked by size.

```
fields @timestamp, source_type, source_arn, source_name, consumed_source_size
| filter source_type = "SPACE"
| stats latest(consumed_source_size) as latest_size, latest(source_name) as name by
  source_arn
| sort latest_size desc
| limit 20
```

All sources detail

Shows the latest state of every source with size, type, name, document count, and last update time.

```
fields @timestamp, source_type, source_name, consumed_source_size,
  consumed_source_doc_count, source_arn
| stats latest(consumed_source_size) as latest_size, latest(source_type) as type,
  latest(source_name) as name, latest(consumed_source_doc_count) as doc_count,
  latest(@timestamp) as last_updated by source_arn
| sort latest_size desc
| limit 100
```

Top users by total size

Shows the top 20 users ranked by total source size across all their sources.

```
fields @timestamp, user_arn, source_arn, consumed_source_size
| parse user_arn "*:user/*/*" as @prefix, @namespace, @username
| stats latest(consumed_source_size) as latest_size, latest(@username) as user by
  source_arn
| stats sum(latest_size) as total_size by user
| sort total_size desc
| limit 20
```

Create a CloudWatch dashboard

You can create a CloudWatch dashboard to visualize index usage metrics. Use the CloudWatch console to create a dashboard and add Logs Insights widgets by using the queries from the preceding section.

1. Open the [CloudWatch console](#).
2. In the navigation pane, choose **Dashboards**.
3. Choose **Create dashboard** and enter a name (for example, IndexUsageMetrics).
4. Add widgets by using the **Logs** widget type. Select your index usage log group and paste the queries from the preceding section.

i Tip

Use the pie chart visualization for the source type breakdown. Use the line chart for size over time. Use the bar chart for top sources. The table visualization works well for the all-sources detail query.

Troubleshooting knowledge bases

When you encounter issues with your Quick knowledge base, you can use this troubleshooting guide to identify and resolve common problems. Knowledge base issues typically involve document synchronization, refresh job failures, or access permissions.

Documents don't appear in your knowledge base

When documents you expect to see don't appear in your knowledge base, several factors might cause this issue.

Common causes:

- **Sync in progress** – Documents might still be processing. Check the refresh status to confirm the refresh is complete.
- **Unsupported file format** – Verify your documents are in a supported format. For supported formats and size limits, see [the section called “File size and content limits”](#).
- **File size too large** – Verify the file is within the size limits. For details, see [the section called “File size and content limits”](#).
- **Insufficient access permissions** – Confirm the knowledge base has proper permissions to access the document source.
- **Document filtering** – Check if filters or exclusion rules prevent certain documents from being indexed.

To troubleshoot:

1. Review the refresh history for error messages related to specific documents that failed to sync.
2. Verify your document formats and file sizes meet requirements.
3. Check your access permissions and connection settings.

Refresh job fails

A refresh job typically fails when there's a configuration error in the knowledge base or data source connection.

Common causes:

- **Permission issues** – The integration lacks sufficient permissions to access the data source.
- **Configuration errors** – Incorrect URLs or data source connection settings.
- **Resource limitations** – Rate limiting from the source system.

To resolve:

1. Check the refresh history details for specific error messages.
2. Verify all connection settings and permissions are correctly configured.
3. Take the recommended action based on the error message.

Refresh job completes with issues

When a refresh job completes with issues, the job processed successfully but encountered problems with some documents.

What this means:

- **Partial success** – Some documents synced successfully while others failed.
- **Document-level errors** – Individual files might have formatting issues, corruption, or access problems.
- **Metadata issues** – Problems with document metadata or associated information.
- **Size or format violations** – Some files might exceed size limits or be in unsupported formats.

To resolve:

1. Review the detailed refresh reports to identify which documents encountered issues.
2. Address the individual document problems.
3. Run another refresh after resolving the issues.

Refresh job succeeds but no documents appear

If a refresh job shows as successful but no documents appear in your knowledge base, check these potential causes.

Common causes:

- **Empty source** – The configured data source location contains no documents.
- **Incorrect path configuration** – The source path or connection settings don't point to the correct location.
- **Document filters** – Inclusion or exclusion criteria might filter out all documents.
- **Read permissions missing** – The job connected successfully but lacked permissions to read the actual documents.

To resolve:

1. Verify your data source configuration points to the correct location.
2. Confirm documents are present in the specified location.
3. Check that appropriate access permissions are configured.
4. Review any document filters that might exclude content.

File format issues during refresh

Quick knowledge bases support specific file formats. Files must meet format, size, and character limit requirements. For the full list of supported formats and limits, see [the section called “File size and content limits”](#).

To resolve format issues:

1. Verify your files meet the format and size requirements.

2. Convert unsupported formats to supported ones.
3. Remove password protection from files.
4. Check that files aren't corrupted.

Access denied errors

Access denied errors typically occur due to authentication or authorization issues.

Common causes:

- **Invalid credentials** – Authentication tokens or passwords might have expired.
- **Insufficient permissions** – The account used in the integration lacks read access to the data source.
- **Network restrictions** – Firewall or security policies block access.
- **SSL/TLS issues** – Certificate problems with secure connections.

To resolve:

1. **Verify authentication credentials** – Confirm that authentication credentials are current and valid. Edit the integration to re-authenticate and generate a new token.
2. **For web crawler data sources** – Verify that secure connections are properly configured and SSL certificates are properly configured and trusted.
3. **Contact your system administrator** – If you continue experiencing access issues, contact your system administrator. They might need to adjust permissions or security settings.

ACL validation errors

If your sync report shows items with status **SKIPPED** and error type **VALIDATION_ERROR** with the message "File has no ACL while crawlACL is true, skipping ingestion," the app registration used by your knowledge base connector is missing the required ACL permissions.

To resolve:

1. Verify the app registration has the correct API permissions for ACL crawling. The required permissions vary by connector. See the permissions section in your connector's setup documentation.

2. Confirm that admin consent has been granted for all required permissions.
3. Re-run a full sync after fixing permissions.

Sync run exceeds maximum runtime

If a sync run takes longer than 14 days, Amazon Quick ends the run with a status of **FAILED** and the message *Maximum sync time exceeded*. This typically happens when a knowledge base is configured to crawl more content than can be processed in a single run.

To resolve:

1. Edit your knowledge base to reduce the volume of content per sync. Apply include or exclude filters, restrict the file types or folders being crawled, or split the content across multiple knowledge bases.
2. For web crawler data sources, consider using the Web Search feature instead if your goal is to chat with large public websites such as `wikipedia.org`.
3. Run the sync again after narrowing the scope.

For more information about this limit, see [the section called "Maximum sync duration"](#).

Connectors

Connectors use secure connections to external services and execute actions based on your authentication level and permissions.

How connectors work

Connectors in Amazon Quick create secure connections between Amazon Quick and external services. When you configure these integrations, you can perform actions based on your authentication level and permissions.

Core components

Connectors

The foundational resources that integrate with external services. For a list of supported connectors, see [Supported integrations](#). For information about setting up AWS built-in service connectors, see [AWS service connectors](#).

Authentication methods

Connectors support multiple authentication methods including Default OAuth app, Custom OAuth app, Service-to-Service OAuth, and API key. For detailed information about each authentication method, see [the section called "Authentication methods"](#).

Implementation types

- **On-demand actions for immediate, user-triggered operations** - Real-time operations that execute immediately when you trigger them. You can initiate actions through chat interfaces, dashboards, or Amazon Q Apps. Examples include creating tickets, sending messages, or querying data.
- **Automated workflows for scheduled or system-triggered tasks** - System-managed operations that execute based on schedules or triggers. They run in the background without user intervention. Examples include data synchronization, report generation, or system maintenance.

Permission models

- **Personal access permissions through 3LO** - You can grant specific permissions to Amazon Quick through Three-Legged OAuth, maintaining control over your service access. Permissions are tied to your identity and credentials in the target service.
- **Service-level permissions for automated workflows** - Applied to automated workflows, these permissions support system-to-system interactions without user involvement. They're configured at the service level and typically use API keys or service account credentials.
- **Entity-level access controls** - Govern access to actions within Amazon Quick, determining which users or groups can execute specific actions. These controls integrate with Amazon Quick's broader permission system for consistent access management across the platform.

Types of actions

Amazon Quick supports two methods of invoking actions, each serving different use cases and authentication models.

On-demand actions

On-demand actions execute immediately when you trigger them. These actions support interactive operations that require real-time response.

Key characteristics:

- User-initiated execution - You trigger actions through natural language in the chat interface.
- Interactive form completion - You fill out forms with required parameters before the action executes.
- Immediate response - Actions execute in real-time and provide instant feedback on success or failure.
- Personal authentication (3LO) - Uses your individual credentials and permissions from the target service.

Common use cases:

- Creating tickets in Jira.
- Sending messages in Slack.
- Updating Salesforce records.
- Retrieving information from SharePoint.

Automated workflows

Automated workflows execute actions on a schedule or in response to specific triggers. These are useful for background and system-level operations.

Key characteristics:

- System-level execution - Actions run automatically without user intervention based on predefined triggers.
- Scheduled or event-triggered - Execute on time-based schedules or in response to specific system events.
- Non-interactive operation - Run in the background without requiring user input or form completion.
- Service-level authentication - Use system credentials rather than individual user authentication.

Common use cases:

- Regular data synchronization.
- Scheduled report generation.
- Automated ticket updates.

- System health checks.

Prerequisites

Before using actions in Amazon Quick, ensure you have the following:

Subscription requirements

For information about subscription requirements for configuring and using connectors, see [Set up integrations in the console](#).

Service requirements

For third-party services (such as Jira or Salesforce), ensure that you have:

- Appropriate permissions in the target services.
- Authentication credentials for each service.

For AWS connectors, you need admin access to the relevant services.

AWS account requirements

- Active AWS account - A valid AWS account with billing enabled and in good standing.
- Appropriate IAM permissions - IAM roles and policies that allow Amazon Quick to access the required AWS services.
- Required service quotas - Sufficient service limits for the AWS services you plan to integrate with your actions.

Tool permissions

Connector owners and administrators can enable or disable individual tools within a connector to control which tools are available to users. Tool permission settings also let connector owners decide which tools require user consent before running, and which tools users can configure for themselves.

Tool permission settings

Connector authors configure tool permissions at the individual action level. The permission setting determines whether you see a confirmation prompt before an action runs. These settings apply across all Amazon Quick clients, including Chat, Agents, Pages, and Flows.

Always Ask

When a connector author sets an action to **Always Ask**, every invocation presents you with a confirmation prompt. You cannot override this behavior through your user-level preferences. The prompt always appears.

Let Users Choose (default)

Let Users Choose is the default setting for all actions. When a connector author sets an action to **Let Users Choose**, your user-level preference determines the initial behavior.

- For write actions, the user-level default is **Always Ask**, so you see a confirmation prompt on every invocation until you change your preference.
- For read actions, the user-level default is **Always Allow**, so you do not see a confirmation prompt unless you change your preference.

When you see a confirmation prompt, you can choose one of the following options.

- **Allow**. The action runs this time. The prompt reappears on the next invocation.
- **Trust**. The action runs and the prompt does not appear for future invocations.
- **Deny**. The action does not run.

Connector-level settings

Connector owners configure tool behavior for all users of a connector.

Connector-level settings

Setting	Behavior
Let Users Choose	You can individually configure how each action runs through your user-level preferences.
Always Ask	You always see a confirmation prompt before the action runs.

Setting	Behavior
Disable	No user can run the action.

User-level settings

You can configure tool behavior for your own connector. These settings apply only to your connection and do not affect other users.

User-level settings

Setting	Behavior
Always Ask	You see a confirmation prompt before the action runs.
Always Allow	The action runs without a confirmation prompt.
Disable	You cannot run the action from your connection.

Default behavior

The following table shows the default settings that Amazon Quick applies to each action type.

Default behavior by action type

Action type	Connector-level default	User-level default	User-changeable
Write actions	Let Users Choose	Always Ask	Yes
Read actions	Let Users Choose	Always Allow	Yes

Supported connector types and available actions

Amazon Quick supports multiple connector types, each with specific actions available:

External service connectors

- **Adobe Marketing Agent** - Work with Adobe Experience Platform marketing data and AI-powered marketing capabilities.
- **Airtable** - Create and update records, manage bases and tables, filter views, and perform relational data operations.
- **Asana** - Manage tasks, workspace operations, and project handling.
- **Atlassian Confluence Cloud** - Create, update, and manage pages, spaces, and other Confluence objects.
- **Atlassian Jira Cloud** - Create issues, update tickets, search projects, manage workflows.
- **BambooHR** - Access employee data, manage time-off requests.
- **Box** - Manage files, folders, and collaborate on documents.
- **Canva** - Create and edit designs, manage templates and assets.
- **Cisco Webex Meetings** - Schedule meetings, retrieve recordings, and manage participants.
- **Cisco Webex Video Messaging** - Manage video updates, retrieve transcripts, and access video metadata in Cisco Webex (Vidcast).
- **Dropbox** - Upload files, manage folder structures, generate sharing links, and control access permissions.
- **Dun & Bradstreet** - Look up company profiles, credit risk data, and business intelligence.
- **GitHub** - Manage repositories, issues, pull requests, and code collaboration.
- **Gmail** - Read emails, search inbox, draft messages, manage labels, and organize threads.
- **Google Analytics** - Query traffic reports, retrieve audience data, monitor conversion events, and access performance metrics.
- **Google Calendar** - Create and update events, check availability, manage invites, and retrieve schedules.
- **Google Chat** - Send messages, manage spaces, and interact with chat threads.
- **Google Docs** - Create and edit documents, read content, and collaborate on text-based files.
- **Google Drive** - Upload files, manage folders, search content, and perform file operations.
- **Google Meet** - Schedule video meetings, list hosted meetings, retrieve transcripts, and update access settings.
- **Google Sheets** - Create and edit spreadsheets, read and write cell ranges, and perform data operations.

- **Google Slides** - Create presentations, edit slides, manage layouts, and add speaker notes.
- **HG Insights** - Access technographic data, query company technology profiles, and retrieve market intelligence.
- **HubSpot** - Manage contacts, deals, marketing campaigns, and CRM data.
- **HuggingFace** - Access AI models, datasets, and machine learning workflows.
- **Intercom** - Manage customer conversations, support tickets, and messaging.
- **Linear** - Create and manage issues, projects, and development workflows.
- **Microsoft OneDrive** - Create, update, delete, and manage OneDrive files and folders.
- **Microsoft OneNote** - Create and edit notebooks, manage sections, and organize notes.
- **Microsoft Outlook** - Send emails, manage calendar events, access contacts.
- **Microsoft SharePoint** - Perform actions on SharePoint lists, items, files, and Excel workbooks.
- **Microsoft Teams** - Send messages, manage channels, schedule meetings, and manage team collaboration.
- **Monday.com** - Manage projects, tasks, and team collaboration workflows.
- **Moody's GenAI Ready Data** - Access Moody's credit ratings, financial research, and risk analytics data.
- **New Relic** - Access observability data, query metrics, and manage monitoring workflows.
- **Notion** - Create and manage pages, databases, and collaborative workspaces.
- **PagerDuty** - Create incidents, manage escalations, update on-call schedules.
- **QuickBooks** - Manage invoices, track expenses, generate financial reports, and handle vendor records.
- **Salesforce** - Create records, update opportunities, search accounts, manage leads.
- **SAP** - Access SAP S/4HANA systems to perform read operations on enterprise data.
- **ServiceNow** - Create incidents, update requests, manage workflows.
- **Shopify** - Manage products, orders, customers, and store operations.
- **Slack** - Send messages, create channels, manage notifications.
- **Smartsheet** - Update sheets, manage project data.
- **Snowflake Cortex Agent** - Query Snowflake data and run AI-powered analytics through Cortex Agent.
- **Visier Agent** - Access workforce analytics, query people data, and generate HR insights.
- **WhatsApp** - Send and manage messages through WhatsApp Business.

- **Zapier** - Trigger workflows across thousands of connected applications.
- **Zendesk** - Create tickets, update cases, search knowledge base.
- **Zoom** - Schedule meetings, add webinar registrants, summarize recorded meetings, and list participants.
- **ZoomInfo** - Look up company and contact data, and query firmographic and intent data.

AWS service connectors

- **Amazon S3** - Upload files, manage buckets, retrieve objects.
- **Amazon Bedrock** - Generate content, analyze data, process requests.
- **Amazon Textract** - Extract text and data from documents.
- **Amazon Comprehend** - Natural language processing and sentiment analysis.
- **Amazon Comprehend Medical** - Medical text analysis and entity extraction.

Connector compatibility matrix

The following table shows which Amazon Quick features each connector type supports:

Connector feature compatibility

Connector	Chat Agents	Flows	Dashboard Visuals	Dashboard Alerts	Automations	Companions
AWS Built-in Services						
AWS Bedrock Agent Runtime	—	—	—	—	✓	—
AWS Bedrock Data Automation Runtime	—	—	—	—	✓	—
AWS Bedrock Runtime	—	—	—	—	✓	—
Amazon Comprehend	—	—	—	—	✓	—

Connector	Chat Agents	Flows	Dashboard Visuals	Dashboard Alerts	Automations	Companions
Amazon Comprehend Medical	—	—	—	—	—	—
Amazon S3	—	—	—	—	✓	—
Amazon Textract	—	—	—	—	✓	—
External Service Connectors						
Adobe Marketing Agent	✓	✓	—	—	—	—
Airtable	✓	✓	✓	—	—	—
Asana	✓	✓	—	—	—	✓
Atlassian Confluence Cloud	✓	✓	—	—	✓	✓
Atlassian Jira Cloud	✓	✓	✓	✓	✓	✓
BambooHR	✓	✓	—	—	—	✓
Box	✓	✓	—	—	—	—
Canva	✓	✓	—	—	—	—
Cisco Webex Meetings	✓	✓	—	—	—	—
Cisco Webex Video Messaging	✓	✓	—	—	—	—
Dropbox	✓	✓	✓	—	—	—
Dun & Bradstreet	✓	✓	—	—	—	—

Connector	Chat Agents	Flows	Dashboard Visuals	Dashboard Alerts	Automations	Companions
GitHub	✓	✓	—	—	—	—
Gmail	✓	✓	✓	—	—	—
Google Analytics	✓	✓	✓	—	—	—
Google Calendar	✓	✓	✓	—	—	—
Google Chat	✓	✓	—	—	—	—
Google Docs	✓	✓	✓	—	—	—
Google Drive	✓	✓	✓	—	—	—
Google Meet	✓	✓	✓	—	—	—
Google Sheets	✓	✓	✓	—	—	—
Google Slides	✓	✓	✓	—	—	—
HG Insights	✓	✓	—	—	—	—
HubSpot	✓	✓	—	—	—	—
HuggingFace	✓	✓	—	—	—	—
Intercom	✓	✓	—	—	—	—
Linear	✓	✓	—	—	—	—
Microsoft OneDrive	✓	✓	✓	✓	✓	✓
Microsoft OneNote	✓	✓	—	—	—	—
Microsoft Outlook	✓	✓	✓	✓	✓	✓

Connector	Chat Agents	Flows	Dashboard Visuals	Dashboard Alerts	Automations	Companions
Microsoft SharePoint	✓	✓	✓	✓	✓	✓
Microsoft Teams	✓	✓	✓	✓	✓	✓
Monday.com	✓	✓	—	—	—	—
Moodys GenAI Ready Data	✓	✓	—	—	—	—
New Relic	✓	✓	—	—	—	—
Notion	✓	✓	—	—	—	—
PagerDuty	✓	✓	—	—	✓	✓
QuickBooks	✓	✓	✓	—	—	—
Salesforce	✓	✓	✓	✓	✓	✓
SAP	✓	—	—	—	✓	✓
ServiceNow	✓	✓	✓	✓	✓	✓
Shopify	✓	✓	—	—	—	—
Slack	✓	✓	✓	—	—	✓
Smartsheet	✓	✓	—	—	—	✓
Snowflake Cortex Agent	✓	✓	—	—	—	—
Visier Agent	✓	✓	—	—	—	—
WhatsApp	✓	✓	—	—	—	—
Zapier	✓	✓	—	—	—	—

Connector	Chat Agents	Flows	Dashboard Visuals	Dashboard Alerts	Automations	Companions
Zendesk	✓	✓	—	—	—	✓
Zoom	✓	✓	✓	—	—	—
ZoomInfo	✓	✓	—	—	—	—

Custom Connector Types

Model Context Protocol (MCP)	✓	✓	—	—	✓	—
OpenAPI	✓	✓	—	—	—	—
REST API	—	—	—	—	✓	—

Authentication Support:

- **Chat Agents and Companions** - Support user authentication (Default OAuth app, Custom OAuth app, Basic)
- **Dashboard Visuals** - Support user authentication (Default OAuth app, Custom OAuth app)
- **Dashboard Alerts** - Support system authentication (Service-to-Service OAuth or API Key)
- **Automations** - Support system authentication (Service-to-Service OAuth)

Connector APIs

Connector APIs let you programmatically create and manage connections between Amazon Quick and external services. These APIs support the action integration functionality that allows users to perform actions in third-party applications directly from Amazon Quick chat interfaces and automated workflows.

What are connector APIs?

Connectors serve as the foundational resources that enable integration with first and third party applications. Through these APIs, you can authenticate to applications, manage permissions, and control which actions are available to users within your Amazon Quick applications.

How connector APIs support action integrations

Connector APIs provide the backend infrastructure for Amazon Quick action integrations. When you create a connector through the API, you establish a secure connection that lets you:

- Execute actions in external services through chat interfaces.
- Perform automated workflows in background processes.
- Integrate third-party services with Amazon Quick applications.
- Manage authentication and permissions for service access.

The APIs handle the complex authentication flows, credential management, and permission controls needed to securely connect Amazon Quick with external services.

Authentication methods

Connector APIs support multiple authentication methods to accommodate different use cases and security requirements:

Managed authentication (3LO)

Three-Legged OAuth provides the simplest setup for personal access to third-party services:

- No initial configuration required.
- User-specific authentication through service provider login.
- Automatic token refresh with 90-day lifecycle.
- Secure credential storage managed by Amazon Quick.

Service-to-service authentication (2LO)

For complex enterprise integrations:

- Supports client credentials OAuth flow.
- Enables system-to-system interactions.
- Requires client ID, client secret, and token URL configuration.
- Suitable for automated workflows requiring sophisticated security.
- OAuth - Dynamic Client Registration (DCR - applicable only for select MCP servers).

API key authentication

Simplified authentication for automated workflows:

- Single token-based authentication.
- Service-level permissions.
- Ideal for background processes and scheduled actions.
- Requires valid API key from target service.

Basic Auth

Basic authentication provides a simple username/password authentication method:

- Uses standard HTTP Basic Authentication headers.
- Credentials are base64 encoded.
- Suitable for services that don't support OAuth or API keys.
- Requires secure HTTPS connection.
- Not recommended for public-facing services.

None

No authentication required:

- Used for public APIs and services.
- No credentials or tokens needed.
- Limited to read-only or public operations.
- Typically used for public data feeds and documentation.
- Should not be used for sensitive operations.

Permissions and access control

Connector APIs implement comprehensive permission controls through Access Control Lists (ACLs):

Resource-level permissions

- **Owner** - Full control including delete and permission management.
- **Contributor** - Can use and modify connector settings.

- **Viewer** - Can view connector details and use enabled actions.

API operations for permission management

- `DescribeActionConnectorPermissions` - Retrieve current permission settings.
- `UpdateActionConnectorPermissions` - Grant or revoke user permissions.

Supported connector categories

Dual-purpose connectors

These connectors support both action integrations and knowledge base creation:

- **Amazon S3** - Use the Admin Console to create Actions for file operations, use the webapp to create knowledge bases from S3 content.
- **Microsoft SharePoint** - Document management actions, content indexing.
- **OneDrive** - File operations, document search capabilities.
- **Confluence** - Content creation actions, knowledge base integration.

Action-only connectors

Specialized for action execution without knowledge base capabilities:

- **Salesforce** - Enterprise CRM integration supporting account and contact operations, custom object CRUD operations, Sales process automation.
- **JIRA** - Issue tracking and project management.
- **Microsoft Outlook** - Send emails, manage calendar events, access contacts.
- **Slack** - Communication and notification workflows.
- **ServiceNow** - IT service management operations.
- **Zendesk** - Create tickets, update cases, search knowledge base.
- **PagerDuty** - Create incidents, manage escalations, update on-call schedules.
- **Asana** - Create actions, update projects, manage team workflows.
- **BambooHR** - Access employee data, manage time-off requests.
- **Smartsheet** - Update sheets, manage project data.
- **FactSet** - Access financial data, generate reports.

- **SAP** - Access SAP systems, execute business functions, and manage enterprise data.

Knowledge base-only connectors

Focused on knowledge base integration without action capabilities:

- **Google Drive** - Document indexing and search.
- **Web Crawler** - Content discovery and indexing.

API lifecycle management

Credential management

- Automatic refresh token handling for OAuth connectors.
- Secure storage of authentication credentials using AWS KMS.
- Support for credential rotation and updates.
- Cross-account access for Amazon S3 connectors.

Connection updates

Use the `UpdateActionConnector` API to:

- Modify authentication credentials.
- Update service configuration parameters.
- Change connector metadata.

Monitoring and troubleshooting

- Track API usage through CloudWatch metrics.
- Monitor connection health and authentication status.
- Implement error handling for common failure scenarios.
- Use validation APIs to diagnose configuration issues.

Rate limiting and quotas

Connector APIs implement standard AWS API rate limiting:

- Standard AWS API throttling applies to all operations.
- Connection validation may have additional limits.
- Action execution rates depend on target service capabilities.
- Implement exponential backoff for retry logic.

Cross-account support

For Amazon S3 connectors, the APIs support cross-account access:

- Specify different AWS account IDs during connector creation.
- Configure appropriate IAM permissions for cross-account access.
- Use AWS KMS for secure credential management across accounts.
- Validate permissions before enabling cross-account connections.

Error handling and troubleshooting

Connector APIs return standard AWS error responses:

Common error types

- `AccessDeniedException` - Insufficient permissions for the operation.
- `InvalidParameterValueException` - One or more parameter values are invalid for the operation.
- Invalid configuration parameters - Service-specific configuration values are incorrect or missing.
- `ResourceNotFoundException` - Connector or resource not found.
- `ThrottlingException` - Rate limit exceeded.
- `ConflictException` - Resource conflict or duplicate names.
- `InternalFailureException` - Internal service error occurred during request processing.
- `ResourceExistsException` - Attempt to create a resource that already exists.
- `InvalidNextTokenException` - The pagination token provided is invalid or expired.
- `AccessTokenNotFoundException` - User needs to authorize the connection (that is, sign-button). This exception is used by UX to ask users for authorization.
- `TokenResponseException` - Action setup is not valid.

Implement proper error handling in your applications to manage these scenarios gracefully and provide meaningful feedback to users.

Using connector APIs with AWS CLI

You can use the AWS CLI to manage connectors programmatically. The following examples demonstrate common operations using generic placeholder values.

Creating a connector

Use the `create-action-connector` command to create a new connector for integrating with external services.

```
aws quicksight create-action-connector \  
  --aws-account-id "123456789012" \  
  --name "MyS3Connector" \  
  --action-connector-id "my-s3-connector-id" \  
  --type "AMAZON_S3" \  
  --authentication-config '{  
    "AuthenticationType": "IAM",  
    "AuthenticationMetadata": {  
      "IamConnectionMetadata": {  
        "RoleArn": "arn:aws:iam::123456789012:role/MyConnectorRole"  
      }  
    }  
  }' \  
  --enabled-actions "CreateBucket" "ListBuckets" \  
  --description "S3 connector for automation workflows" \  
  --region "us-east-1"
```

Listing connectors

Use the `list-action-connectors` command to retrieve all connectors in your account.

```
aws quicksight list-action-connectors \  
  --aws-account-id "123456789012" \  
  --max-results 10 \  
  --region "us-east-1"
```

Describing a connector

Use the `describe-action-connector` command to get detailed information about a specific connector.

```
aws quicksight describe-action-connector \  
  --aws-account-id "123456789012" \  
  --action-connector-id "my-s3-connector-id" \  
  --region "us-east-1"
```

Updating a connector

Use the `update-action-connector` command to modify an existing connector's configuration.

```
aws quicksight update-action-connector \  
  --aws-account-id "123456789012" \  
  --action-connector-id "my-s3-connector-id" \  
  --name "UpdatedS3Connector" \  
  --authentication-config '{  
    "AuthenticationType": "IAM",  
    "AuthenticationMetadata": {  
      "IamConnectionMetadata": {  
        "RoleArn": "arn:aws:iam::123456789012:role/UpdatedConnectorRole"  
      }  
    }  
  }' \  
  --enabled-actions "CreateBucket" "ListBuckets" "DeleteBucket" \  
  --region "us-east-1"
```

Searching connectors

Use the `search-action-connectors` command to find connectors based on specific criteria.

```
aws quicksight search-action-connectors \  
  --aws-account-id "123456789012" \  
  --max-results 5 \  
  --filters '[{
```



```
"Name": "ACTION_CONNECTOR_NAME",
"Operator": "StringLike",
"Value": "S3"
}]' \
--region "us-east-1"
```

Managing connector permissions

Use the `update-action-connector-permissions` command to grant or revoke permissions for a connector.

```
aws quicksight update-action-connector-permissions \
--aws-account-id "123456789012" \
--action-connector-id "my-s3-connector-id" \
--grant-permissions '[{
  "Actions": [
    "quicksight:DescribeActionConnector",
    "quicksight:UpdateActionConnector",
    "quicksight:DeleteActionConnector"
  ],
  "Principal": "arn:aws:quicksight:us-east-1:123456789012:user/default/myuser"
}]' \
--region "us-east-1"
```

Viewing connector permissions

Use the `describe-action-connector-permissions` command to view current permissions for a connector.

```
aws quicksight describe-action-connector-permissions \
--aws-account-id "123456789012" \
--action-connector-id "my-s3-connector-id" \
--region "us-east-1"
```

Deleting a connector

Use the `delete-action-connector` command to remove a connector from your account.

```
aws quicksight delete-action-connector \  
  --aws-account-id "123456789012" \  
  --action-connector-id "my-s3-connector-id" \  
  --region "us-east-1"
```

Next steps

After understanding connector APIs, you can:

- Review the complete API reference documentation for detailed parameter specifications.
- Explore specific connector setup guides for your target services.
- Implement authentication flows appropriate for your use case.
- Set up monitoring and error handling for production deployments.
- Configure permissions and access controls for your organization.

Authentication methods

Amazon Quick supports multiple authentication methods for connectors. The available methods depend on the specific connector. During setup, you choose your authentication method from the **OAuth Configuration** options or from the connector-specific authentication settings.

Default OAuth app

Default OAuth app (also known as managed authentication or 3LO) is the recommended authentication method for personal access to third-party services. With this method, Amazon Quick manages the OAuth flow and no additional credentials are required from your organization.

Key features:

- No additional credentials or configuration required.
- User-specific authentication through provider sign-in.
- Secure credential storage managed by Amazon Quick.
- Automatic token refresh.
- 90-day refresh token lifecycle.

Setup process:

1. Choose **Default OAuth app** as your OAuth configuration.
2. Complete the service provider sign-in flow.
3. Grant the requested permissions.
4. Confirm the connection.

Note

Not all connectors support Default OAuth app. Check the connector-specific documentation for available authentication methods.

Custom OAuth app

Custom OAuth app authentication is for organizations that require specific control over the OAuth application configuration. Like Default OAuth app, this method uses a three-legged OAuth (3LO) flow where users sign in directly with the service provider. The difference is that you provide your own OAuth credentials instead of using the Amazon Quick managed application.

Required information:

- Client ID
- Client Secret
- Domain URL
- Authorization URL
- Token URL
- Redirect URL

Setup process:

1. Create an OAuth application in your service provider's developer console.
2. Choose **Custom OAuth app** as your OAuth configuration.
3. Enter the credentials from your OAuth application.
4. Complete the sign-in flow and validate the connection.

Service-to-Service OAuth

Service-to-Service OAuth uses client credentials for server-to-server authentication without user interaction. This method is suitable for automated workflows and shared connectors where actions run under a service account.

Required information:

- Client ID
- Client Secret
- Domain URL
- Token URL
- Service-specific parameters (varies by connector)

API key authentication

Some connectors support API key authentication for service-level access. This method uses a single token for authentication and is common for connectors that don't support OAuth.

Required information:

- Valid API key from the service provider
- Base URL or domain
- Service-specific parameters (such as email or account ID)

Action execution methods

Amazon Quick provides multiple ways to execute actions, accommodating different use cases and interaction preferences.

Chat interface

You can execute implicit actions in the Amazon Quick chat.

Implicit actions

Amazon Quick also supports implicit action execution through natural conversation with agents. Using advanced natural language processing, the system can identify when your conversation

indicates the need for specific actions. Conversations are analyzed to determine which actions are required to fulfill your request.

A single request might require multiple actions to complete. When this happens, the system handles these actions sequentially, guiding you through each step. For each identified action, the system presents the appropriate form for you to complete. After each action completes, you receive a confirmation before moving on to the next action in the sequence.

For example, if you ask "Create a Jira ticket for this issue and notify the team in Slack," the system would:

1. First present the Jira ticket creation form.
2. After completing the ticket creation, show the Slack message form.
3. Complete both actions in sequence.

Throughout the process, you can track your progress through multiple actions. When all actions complete, the system provides a comprehensive summary showing all executed actions and their outcomes. You can access related documentation if needed and review any error states that may have occurred during the process.

Monitoring and maintenance

Monitoring your connectors helps ensure reliable performance and identify issues before they impact users. Regular monitoring allows you to track usage patterns, optimize performance, and maintain healthy connections to external services.

Performance monitoring

You can assess connector performance using the following metrics and analytics.

CloudWatch metrics

- Action execution success rates - Track the percentage of successful action executions to identify reliability issues.
- Response times - Monitor how long actions take to complete and identify performance bottlenecks.
- Error frequencies - Track error patterns to identify common failure points and areas for improvement.

- API quota usage - Monitor usage against service limits to prevent throttling and plan for capacity.

Usage analytics

The following usage analytics are collected for connectors:

- Active users - Track how many users are actively using connectors to understand adoption and usage patterns.
- Popular actions - Identify which actions are used most frequently to prioritize optimization efforts.
- Execution patterns - Analyze when and how often actions are executed to optimize resource allocation.
- Error trends - Monitor error patterns over time to identify systemic issues and improvement opportunities.

Connection health

You can assess connector health using the following connection health tools:

Status monitoring

- Connection state - Monitor whether connectors are actively connected and functioning properly.
- Authentication validity - Track the status of authentication tokens and credentials to prevent access failures.
- Token expiration tracking - Monitor when authentication tokens will expire and need renewal.
- Service availability - Track the availability and response status of connected external services.

Automated maintenance

- Token refresh handling.
- Connection recovery.
- Error retry logic.
- Performance optimization.

CloudWatch metrics reference

Available CloudWatch metrics

Metric	Description	Unit
ActionSuccess	Successful executions	Count
ActionLatency	Execution time	Milliseconds
AuthFailures	Failed authentications	Count
APIThrottling	API throttling events	Count

Best practices

Following best practices for connectors helps ensure secure, reliable, and efficient operations. These practices help you maintain optimal performance, protect sensitive data, and minimize operational issues.

Security

Authentication management

- Regular credential rotation - Update API keys and OAuth tokens on a scheduled basis to maintain security.
- Periodic permission reviews - Audit user and service permissions quarterly to ensure least-privilege access.
- Token lifecycle monitoring - Track token expiration dates and set up alerts before credentials expire.
- Access audit logging - Enable comprehensive logging to track who accessed which services and when.

Access control

- Implement least-privilege access - Grant only the minimum permissions necessary for each action to function properly.

- Regular permission audits - Review and validate that current permissions align with actual usage patterns and business needs.
- Document access patterns - Maintain clear documentation of who has access to which connectors and why.
- Monitor usage anomalies - Set up alerts for unusual access patterns that might indicate security issues.

Performance

Action configuration

- Optimize form defaults - Pre-populate commonly used values to reduce user input time and errors.
- Configure appropriate timeouts - Set realistic timeout values based on typical response times for each service.
- Set up error handling - Implement robust error handling with clear user messages and retry logic where appropriate.
- Document dependencies - Clearly document any prerequisites or dependencies between different actions.

Resource management

- Monitor API quotas.
- Track usage patterns.
- Optimize refresh schedules.
- Regular cleanup of unused connectors.

Maintenance

Regular actions

- Review connector status.
- Update configurations.
- Validate connections.

- Document changes.

Troubleshooting

- Monitor error patterns.
- Review CloudWatch logs.
- Track resolution times.
- Document solutions.

Troubleshooting

When connectors encounter issues, systematic troubleshooting helps you quickly identify and resolve problems. This guidance covers common issues and their solutions to minimize downtime and restore functionality.

Common issues and solutions

Authentication problems

Token expiration

Symptom: "Authentication token expired" error
Resolution:

1. Choose "Reconnect" in console.
2. Complete authentication flow.
3. Retry action.

Permission errors

Symptom: "Insufficient permissions" message
Resolution:

1. Verify service permissions.
2. Check connector configuration.
3. Review action requirements.

Connection failures

Symptom: "Unable to connect to service" error

Resolution:

1. Verify service availability.
2. Check network connectivity.
3. Validate credentials.
4. Review service quotas.

Action-specific issues

Form submission failures

Validation errors

- Check required fields.
- Verify data formats.
- Review field limitations.
- Check for special characters.

Timeout issues

- Reduce form complexity.
- Check network latency.
- Review service response times.
- Consider breaking into multiple actions.

Sync and performance issues

Slow response times

Resolution:

1. Check API rate limits.
2. Review concurrent executions.

- 3. Monitor service health.
- 4. Optimize action configuration.

Failed executions

Resolution:

- 1. Review CloudWatch logs.
- 2. Check error messages.
- 3. Verify service status.
- 4. Test connection health.

Common error messages

Error codes and resolutions

Error code	Description	Resolution
AUTH_001	Authentication failed	Verify credentials and retry
CONN_002	Connection timeout	Check network and service status
PERM_003	Insufficient permissions	Review required permissions
TOKEN_004	Token expired	Reinitiate authentication

Integration-specific guides

Use these guides to configure integrations with specific applications. Each guide includes setup instructions for both action and data ingestion capabilities when the integration supports them.

Adobe Marketing Agent integration

With the Adobe Marketing Agent connector, you can work with Adobe Experience Platform marketing data and AI-powered marketing capabilities directly in Amazon Quick through natural language.

Amazon Quick supports multiple authentication methods for Adobe Marketing Agent. Choose the method that best fits your organization's security requirements.

- **Default OAuth app** – Uses an AWS-managed OAuth application. No additional credentials are needed. Users authenticate directly with their Adobe account.
- **Custom OAuth app** – Uses a customer-managed application registered in Adobe Developer Console. This option gives your organization full control over the OAuth configuration.
- **Service-to-Service OAuth** – Uses client credentials for server-to-server authentication without user interaction. Suitable for automated workflows.
- **API Key** – Uses an API key for authentication. Suitable for service-account access patterns.

For more information about the authentication methods that Amazon Quick supports, see [the section called “Authentication methods”](#).

Before you begin

Make sure that you have the following before you set up the integration.

- An active Adobe account with access to Adobe Marketing Agent.
- For **Custom OAuth app** or **Service-to-Service OAuth**: Access to [Adobe Developer Console](#) on the Adobe website to create an app and obtain client credentials.
- For **API Key**: An Adobe-issued API key with the required marketing data access scopes.
- For Amazon Quick subscription requirements, see [the section called “Set up integrations in the console”](#).

Configuring Adobe

If you are using **Default OAuth app** authentication, skip this section and proceed to [the section called “Setting up the connector in Amazon Quick”](#).

For Custom OAuth app or Service-to-Service OAuth authentication, register an app in [Adobe Developer Console](#) on the Adobe website and add the Amazon Quick callback URL `https://{region}.quicksight.aws.amazon.com/sn/oauthcallback` as a redirect URI. Replace `{region}` with your AWS Region (for example, `us-east-1`). Record the Client ID and Client Secret — you need them when you configure Amazon Quick. For API Key authentication, generate an API key from your Adobe account.

Setting up the connector in Amazon Quick

Connect from the Available tab

If you want to use Default OAuth app authentication, you can connect directly from the **Available** tab without additional configuration.

1. In the Amazon Quick console, choose **Connectors**.
2. On the **Available** tab, find **Adobe Marketing Agent** and choose **Connect**.
3. Complete the Adobe sign-in flow and grant the requested permissions.

To configure a connector with one of the other authentication methods, use the **Create for your team** tab as described below.

Create from the Create for your team tab

1. In the Amazon Quick console, choose **Connectors**.
2. Choose the **Create for your team** tab.
3. Find and choose **Adobe Marketing Agent**.
4. Enter a **Name** for your connector. Optionally, choose **+ Add Description** to add a description.
5. For **Connection type**, choose **Public network**.
6. For **OAuth Configuration**, choose one of the following authentication methods and configure the required fields.
 - a. For **Default OAuth app**:

No additional credentials are needed. Choose **Next** to continue.

- b. For **Custom OAuth app**, configure the following fields:
 - **Client ID** – The client ID from your Adobe Developer Console app.
 - **Public OAuth client** (Optional) – Select this option if your Adobe Developer Console app is configured as a public client (no client secret).
 - **Client secret** – The client secret from your Adobe Developer Console app.
 - **Token URL** – The token endpoint. Default: `https://aep-ai-ama.adobe.io/token`
 - **Authorization URL** – The authorization endpoint. Default: `https://aep-ai-ama.adobe.io/authorize`

- **Redirect URL** – Pre-filled with the Amazon Quick callback URL.
- c. For **Service-to-Service OAuth**, configure the following fields:
 - **Client ID** – The client ID from your Adobe Developer Console Server-to-Server credential.
 - **Client secret** – The client secret from your Adobe Developer Console Server-to-Server credential.
 - **Token URL** – The token endpoint. Default: `https://aep-ai-ama.adobe.io/token`
 - d. For **API Key**, configure the following fields:
 - **API Key** – The Adobe API key.
 - **Email (Optional)** – The email address associated with the API key.
7. Choose **Next**.
 8. If you chose **Default OAuth app** or **Custom OAuth app**, an Adobe authorization window opens. Review the requested permissions and choose **Allow**.
 9. On the **Review** page, review the available actions for the connector. Choose **Next**.
 10. On the **Publish** page, choose who can access the connector. You can enable access for everyone in your organization or search for specific teams or groups.
 11. Choose **Publish**.

Available actions

After you set up the connector, the actions exposed by Adobe Marketing Agent are available. To see the current set of actions for your connector, go to the connector's **Available actions** view in the Amazon Quick console.

Managing and troubleshooting

To edit, share, or delete your connector, see [the section called “Managing existing integrations”](#).

Authentication issues

- **Sign-in fails (Default OAuth app or Custom OAuth app)** – Verify that your Adobe account is active and that you can sign in to Adobe directly. For Custom OAuth app, confirm that the redirect URI in your Adobe Developer Console app matches the Amazon Quick callback URL.

- **Invalid client credentials (Custom OAuth app or Service-to-Service OAuth)** – Verify that the Client ID and Client secret match the values in your Adobe Developer Console app.
- **API Key authentication fails** – Verify that the API key has not been revoked and that it has the required scopes for the operations you want to perform.

Airtable integration

With the Airtable connector, you can access the Airtable platform directly in Amazon Quick through natural language. You can create and update records, manage bases and tables, filter views, and perform relational data operations without leaving Amazon Quick.

Amazon Quick supports multiple authentication methods for Airtable. Choose the method that best fits your organization's security requirements.

- **Default OAuth app** – Uses an AWS-managed OAuth application. No additional credentials are needed. Users authenticate directly with their Airtable account.
- **Custom OAuth app** – Uses a customer-managed OAuth application registered in Airtable. This option gives your organization full control over the OAuth configuration.
- **API Key** – Uses an Airtable personal access token for authentication. Suitable for individual use or testing.

For more information about the authentication methods that Amazon Quick supports, see [the section called “Authentication methods”](#).

Before you begin

Make sure that you have the following before you set up the integration.

- An active Airtable account with access to the bases and data that you want to use.
- For **Custom OAuth app**: Access to the [Airtable Builder Hub](#) on the Airtable website to register an OAuth integration.
- For **API Key**: A personal access token from the [Airtable Builder Hub](#) on the Airtable website.
- For Amazon Quick subscription requirements, see [the section called “Set up integrations in the console”](#).

Configuring Airtable

If you are using **Default OAuth app** authentication, skip this section and proceed to [the section called "Setting up the connector in Amazon Quick"](#).

For Custom OAuth app or API Key authentication, complete the applicable steps in Airtable before configuring Amazon Quick.

Register an OAuth integration (Custom OAuth app)

Create an OAuth integration in the Airtable Builder Hub to obtain the client credentials that you need for Amazon Quick.

1. Sign in to the [Airtable Builder Hub](#) on the Airtable website and choose **OAuth integrations**.
2. Choose **Register new OAuth integration**.
3. Enter a **Name** for your integration.
4. For **OAuth redirect URLs**, enter the redirect URL from Amazon Quick:
`https://{region}.quicksight.aws.amazon.com/sn/oauthcallback`
5. Choose **Register integration**.
6. After registration, generate a client secret by choosing **Regenerate client secret**.
7. Record the following values. You need them when you configure Amazon Quick.
 - **Client ID**
 - **Client secret**
8. Under **Scopes**, select the permissions that your integration requires. For the full list of recommended scopes, see [the section called "Recommended scopes"](#).

Create a personal access token (API Key)

Create a personal access token in the Airtable Builder Hub for API Key authentication.

1. Sign in to the [Airtable Builder Hub](#) on the Airtable website and choose **Personal access tokens**.
2. Choose **Create token**.
3. Configure the token name, scopes, and base access as needed for your use case.
4. Copy and securely store the generated token. You need it when you configure Amazon Quick.

Recommended scopes

When you register an OAuth integration or create a personal access token, configure the following scopes based on the actions that you want to use.

Airtable recommended scopes

Scope	Description
<code>data.records:read</code>	Read data in records.
<code>data.records:write</code>	Create, edit, and delete records.
<code>data.recordComments:read</code>	Read comments in records.
<code>data.recordComments:write</code>	Create, edit, and delete record comments.
<code>schema.bases:read</code>	Read the structure of bases, including table names and field types.
<code>schema.bases:write</code>	Edit the structure of bases, including table names and field types.
<code>user.email:read</code>	Read the email address of the authenticated user.

Setting up the connector in Amazon Quick

Connect from the Available tab

If you want to use Default OAuth app authentication, you can connect directly from the **Available** tab without additional configuration.

1. In the Amazon Quick console, choose **Connectors**.
2. On the **Available** tab, find **Airtable** and choose **Connect**.
3. Complete the Airtable sign-in flow and grant the requested permissions.

To configure a connector with Custom OAuth app or API Key instead, use the **Create for your team** tab as described below.

Create from the Create for your team tab

After you complete any required Airtable configuration, create the connector in Amazon Quick.

1. In the Amazon Quick console, choose **Connectors**.
2. Choose the **Create for your team** tab.
3. Find and choose **Airtable**.

Note

If an Airtable connector already exists, a dialog appears with your existing connectors. To use an existing connector, choose it. To create a new one, choose **No, create new**.

4. Enter a **Name** for your connector. Optionally, choose **+ Add Description** to add a description.
5. For **Connection type**, choose **Public network**.
6. For **OAuth Configuration**, choose one of the following authentication methods and configure the required fields.

a. For **Default OAuth app**:

No additional credentials are needed. Choose **Next** to continue.

b. For **Custom OAuth app**, configure the following fields:

- **Base URL** (Optional) – The Airtable API base URL. Example: `https://api.airtable.com`
- **Client ID** – The client ID from your Airtable OAuth integration.
- **Client secret** – The client secret from your Airtable OAuth integration.
- **Token URL** – The token endpoint. Example: `https://airtable.com/oauth2/v1/token`
- **Authorization URL** – The authorization endpoint. Example: `https://airtable.com/oauth2/v1/authorize`
- **Redirect URL** – Pre-filled with the Amazon Quick callback URL.

c. For **API Key**, configure the following fields:

- **Base URL** (Optional) – The Airtable API base URL. Example: `https://api.airtable.com`
- **API Key** – Your Airtable personal access token.

- **Email (Optional)** – The email address associated with your Airtable account.

7. Choose **Next**.
8. If you chose **Default OAuth app** or **Custom OAuth app**, an Airtable authorization window opens. Review the requested permissions, choose the bases and workspaces to grant access to, and choose **Grant access**.
9. On the **Review** page, review the available actions for the connector. Choose **Next**.
10. On the **Publish** page, choose who can access the connector. You can enable access for everyone in your organization or search for specific teams or groups.
11. Choose **Publish**.

Available actions

After you set up the connector, the following actions are available.

Airtable available actions

Category	Action	Description
Records	List Records	Lists records in a table with optional filtering and sorting.
Records	Get Record	Retrieves a specific record by ID.
Records	Batch Create Record	Creates one or more new records in a table.
Records	Update Record	Updates fields on an existing record.
Records	Batch Update Record	Updates multiple records in a single operation using PATCH.
Records	Batch Put Record	Upserts multiple records in a single operation using PUT.
Comments	List Comments	Lists comments on a record.
Comments	Create Comment	Adds a comment to a record.
Comments	Update Comment	Edits an existing comment on a record.

Category	Action	Description
Schema	List Bases	Lists all bases accessible to the authenticated user.
Schema	Get Base Schema	Retrieves the schema of a base, including tables and field definitions.
Schema	Create Table	Creates a new table in a base.
Schema	Update Table	Updates table properties such as name or description.
Schema	Create Field	Adds a new field to a table.
Schema	Update Field	Updates field properties such as name or type.
Attachments	Create Attachment	Uploads an attachment to a record.
Users	Get User Info	Retrieves information about the authenticated user.

Note

The actions that you can use depend on the bases and workspaces accessible to the authenticated user.

Managing and troubleshooting

To edit, share, or delete your connector, see [the section called “Managing existing integrations”](#).

Authentication issues

- **Sign-in fails (Default OAuth app or Custom OAuth app)** – Verify that your Airtable account is active and that you can sign in to airtable.com on the Airtable website directly. For Custom OAuth app, confirm that the redirect URL in your Airtable OAuth integration matches the Amazon Quick callback URL.

- **Invalid client credentials (Custom OAuth app)** – Verify that the Client ID and Client secret match the values in your Airtable OAuth integration. If you regenerated the client secret, update the value in Amazon Quick.
- **API Key rejected** – Verify that your personal access token is active and has the required scopes. Tokens can be revoked or expired in the Airtable Builder Hub.
- **Insufficient permissions** – Verify that the scopes configured for your authentication method include the permissions required for the actions that you want to use. See [the section called “Recommended scopes”](#).

Amazon S3 integration

With Amazon S3 integration in Amazon Quick, you can create knowledge bases from documents stored in S3 buckets. This integration supports data ingestion capabilities for indexing and searching S3 content.

Note

This guide covers Amazon S3 data ingestion integration for knowledge base creation. For Amazon S3 connectors that perform Amazon S3 operations such as uploading, downloading, and deleting files, see [the section called “AWS service connectors”](#). Amazon S3 actions are only supported for Quick Automate.

What you can do

Amazon S3 users can ask questions about content stored in their Amazon S3 buckets. For example, users can inquire about key findings from documents, search for specific information across multiple file types, or analyze data patterns.

The integration enables users to quickly access and understand information from their Amazon S3 content, regardless of file location or type. It also provides contextual details such as modification dates and file metadata, contributing to more efficient information discovery and better-informed decision making.

Before you begin

Before you set up Amazon S3 integration, make sure you have the following:

- AWS account with Amazon S3 access.
- Amazon S3 bucket with documents to index.
- Amazon Quick Enterprise subscription.
- Necessary permissions to create Amazon S3 integrations.
- Your administrator must grant Amazon Quick access to the Amazon S3 buckets you want to use. For more information, see [the section called “Grant Amazon Quick access to Amazon S3 buckets”](#).

Note

Cross-account Amazon S3 access is only supported within the same AWS region.

Administrator setup

Before users can create Amazon S3 integrations and knowledge bases, an Amazon Quick administrator must complete the following setup tasks.

Grant Amazon Quick access to Amazon S3 buckets

Grant Amazon Quick access to the Amazon S3 buckets your organization needs. This applies whether the buckets are in the same AWS account or a different account.

1. In the Amazon Quick admin console, under **Permissions**, choose **AWS resources**.
2. Under **Allow access and autodiscovery for these resources**, select the **Amazon S3** checkbox.
3. Choose **Select S3 buckets**.
4. In the **Select Amazon S3 buckets** dialog, choose the tab that matches your bucket location:
 - **S3 Buckets Linked To Quick Account** – Select the buckets from the list that you want Amazon Quick to access. Selected buckets have read-only permissions by default.
 - **S3 Buckets You Can Access Across AWS** – For cross-account buckets, make sure the account owner has authorized your account. Choose **Use a different bucket**, enter the bucket name, and choose **Add S3 bucket**.
5. (Optional) For cross-account buckets, select **Restrict bucket access to knowledge base creator** to limit access so that only the user who creates the knowledge base can use the bucket.
6. Choose **Finish**.

The selected buckets are now accessible to users during knowledge base creation.

Prepare IAM role and policy setup

Amazon S3 integration uses AWS authentication to access your Amazon S3 buckets. Prepare your IAM role and policy configuration before users set up the integration.

Required IAM permissions

Make sure your AWS account has the following minimum permissions for the Amazon S3 bucket:

- `s3:GetObject` – Read objects from the bucket.
- `s3:ListBucket` – List bucket contents.
- `s3:GetBucketLocation` – Get bucket region information.
- `s3:GetObjectVersion` – Get object versions.
- `s3:ListBucketVersions` – List bucket versions.

Configure Amazon S3 bucket permissions for cross-account access

If you're accessing Amazon S3 buckets in a different AWS account, you must configure IAM policies in the source AWS account.

To configure Amazon S3 bucket permissions for cross-account access

1. Sign in to the AWS Management Console for the account that contains the Amazon S3 bucket.
2. Open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
3. Choose the bucket that you want to grant access to.
4. Choose **Permissions**, and then choose **Bucket Policy**.
5. Add a bucket policy with the following elements:
 - `Version` – Set to "2012-10-17"
 - `Statement` – Array containing policy statements with:
 - `Sid` – "AllowQuickSuiteS3Access"
 - `Effect` – "Allow"
 - `Principal` – AWS ARN for the Amazon Quick service role in your account.
For example, the principal should look like this: `"Principal": { "AWS":`

```
"arn:aws:iam::<quick_account_id>:role/service-role/aws-quicksight-
service-role-v0" }
```

- Action – Array of Amazon S3 permissions: s3:GetObject, s3:ListBucket, s3:GetBucketLocation, s3:GetObjectVersion, s3:ListBucketVersions
- Resource – "*" (applies to the current key), the Amazon S3 bucket path should look like the following: "Resource": ["arn:aws:s3:::bucket_name"]

6. Choose **Save changes**.

Configure KMS key permissions (if your bucket uses encryption)

If your Amazon S3 bucket uses AWS KMS encryption, complete the following steps.

To configure KMS key permissions

1. Open the AWS Key Management Service (AWS KMS) console at <https://console.aws.amazon.com/kms>.
2. Choose the KMS key that is used to encrypt your Amazon S3 bucket.
3. Choose **Key policy**, and then choose **Edit**.
4. Add a statement to the key policy with the following structural elements:
 - Sid – "AllowQuickSuiteKMSAccess"
 - Effect – "Allow"
 - Principal – AWS ARN for the Amazon Quick service role in your account.
For example, the principal should look like this: "Principal": { "AWS": "arn:aws:iam::<quick_account_id>:role/service-role/aws-quicksight-service-role-v0" }
 - Action – Array of KMS permissions: kms:Decrypt, kms:DescribeKey
 - Resource – "*" (applies to the current key), the Amazon S3 bucket path should look like the following: "Resource": ["arn:aws:s3:::bucket_name"]
5. Choose **Save changes**.
6. Wait 2-3 minutes for the policy changes to propagate.

Configure VPC access for Amazon S3 Connector in Amazon Quick

VPC permissions ensure Amazon Quick can only access your Amazon S3 bucket through secure VPC or VPC endpoint connections.

Required policy change

Add this statement to your bucket access policy to allow Amazon Quick to access your bucket through VPC endpoints:

```
{
  "Sid": "Allow-Quick-access"
  ,
  "Principal": "arn:aws:iam::Quick Account:role/service-role/aws-quicksight-service-
role-v0",
  "Action": "s3:*",
  "Effect": "Allow",
  "Resource": [
    "arn:aws:s3:::amzn-s3-demo-bucket",
    "arn:aws:s3:::amzn-s3-demo-bucket/*"
  ],
  "Condition": {
    "Null": {
      "aws:SourceVpce": "false"
    }
  }
}
```

- Replace `amzn-s3-demo-bucket` with your bucket name.
- Replace `Quick Account` with your Amazon Quick account.

The `"aws:SourceVpce": "false"` condition ensures Amazon Quick can only access your bucket through VPC endpoints, maintaining your security requirements.

Deny policies

If your bucket has a policy that restricts traffic to a specific VPC or VPC endpoint via Deny Policy, you must reverse this policy because deny policies take precedence over allow policies.

For example:

```
{
  "Version": "2012-10-17"
  ,
```

```

    "Id": "Policy1415115909152",
    "Statement": [
      {
        "Sid": "Access-to-specific-VPCE-only",
        "Principal": "*",
        "Action": "s3:*",
        "Effect": "Deny",
        "Resource": ["arn:aws:s3:::amzn-s3-demo-bucket",
                     "arn:aws:s3:::amzn-s3-demo-bucket/*"],
        "Condition": {
          "StringNotEquals": {
            "aws:SourceVpce": "vpce-0abcdef1234567890"
          }
        }
      }
    ]
  }
}

```

Should be reversed into:

```

{
  "Version": "2012-10-17",
  "Id": "Policy1415115909152",
  "Statement": [
    {
      "Sid": "Access-to-specific-VPCE-only",
      "Principal": "*",
      "Action": "s3:*",
      "Effect": "Allow",
      "Resource": ["arn:aws:s3:::amzn-s3-demo-bucket",
                   "arn:aws:s3:::amzn-s3-demo-bucket/*"],
      "Condition": {
        "StringEquals": {
          "aws:SourceVpce": "vpce-0abcdef1234567890"
        }
      }
    }
  ]
}

```

Best practices

Restrict access to your Amazon Quick role

Access policies should enforce that the caller is your Amazon Quick role ARN or, at minimum, your Amazon Quick account. This ensures that despite allowing VPC traffic, calls come only from expected sources.

Security recommendations

- Restrict policies to your Amazon Quick role for most secure traffic
- Review your bucket policies regularly to ensure they follow the principle of least privilege

Restrict Amazon S3 bucket access with IAM policy assignments

You can control which Amazon S3 buckets your Amazon Quick users can use to create knowledge bases by creating IAM policies and assigning them to specific users, groups, or all users through Amazon Quick IAM policy assignments. This allows you to restrict who can create knowledge bases against specific buckets, including ACL-aware knowledge bases.

Note

IAM policies assigned through Amazon Quick take precedence over AWS resource-level policies. To ensure your access requirements are met, configure your IAM policies appropriately.

For example, you can assign a restrictive policy to specific users who need access to ACL-aware buckets, while assigning a broader policy to all users for non-ACL buckets.

Step 1: Create an Amazon S3 access policy in IAM

Create an IAM policy in the AWS IAM console that defines which Amazon S3 buckets users can access for knowledge base creation. The following example policy grants access to two specific buckets:

```
{
  "Version": "2012-10-17"
,
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "s3:ListAllMyBuckets",
      "Resource": "arn:aws:s3:::*"
    }
  ],
}
```

```

{
  "Action": [
    "s3:ListBucket"
  ],
  "Effect": "Allow",
  "Resource": [
    "arn:aws:s3:::amzn-s3-demo-bucket-1",
    "arn:aws:s3:::amzn-s3-demo-bucket-2"
  ]
},
{
  "Action": [
    "s3:GetObject",
    "s3:GetObjectVersion"
  ],
  "Effect": "Allow",
  "Resource": [
    "arn:aws:s3:::amzn-s3-demo-bucket-1/*",
    "arn:aws:s3:::amzn-s3-demo-bucket-2/*"
  ]
},
{
  "Action": [
    "s3:ListBucketMultipartUploads",
    "s3:GetBucketLocation"
  ],
  "Effect": "Allow",
  "Resource": [
    "arn:aws:s3:::amzn-s3-demo-bucket-1",
    "arn:aws:s3:::amzn-s3-demo-bucket-2"
  ]
},
{
  "Action": [
    "s3:PutObject",
    "s3:AbortMultipartUpload",
    "s3:ListMultipartUploadParts"
  ],
  "Effect": "Allow",
  "Resource": [
    "arn:aws:s3:::amzn-s3-demo-bucket-1/*",
    "arn:aws:s3:::amzn-s3-demo-bucket-2/*"
  ]
}

```

```
]
}
```

Replace `amzn-s3-demo-bucket-1` and `amzn-s3-demo-bucket-2` with the names of the Amazon S3 buckets you want to grant access to.

Step 2: Assign the policy in Amazon Quick

After creating the IAM policy, assign it to Amazon Quick users or groups.

1. In the Amazon Quick admin console, under **Permissions**, choose **IAM policy assignments**.
2. Choose **Add new assignment**.
3. Enter a name for the assignment.
4. On the **Select an IAM policy** page, search for and select the IAM policy you created in Step 1. Choose **Next**.
5. On the **Assign users and groups** page, choose one of the following:
 - Select **Assign to all users and groups** to apply the policy to all current and future users.
 - Search for and select specific users or groups to assign the policy to.

Choose **Next**.

6. On the **Review and enable changes** page, verify your assignment details and choose **Save and enable**.

Users who are not explicitly granted access through an IAM policy assignment will not be able to access the restricted Amazon S3 buckets for creating integrations or knowledge bases.

Document-level ACLs

You can enable access control lists (ACLs) at the Amazon S3 knowledge base level using one of two configuration methods, each optimized for different use cases.

Important

Document-level ACL configuration is permanent. You cannot enable ACLs on knowledge bases created without ACL support, and you cannot disable ACLs once enabled. To change ACL configuration, create a new knowledge base with your desired setting from the start.

Note

For ACL-enabled knowledge bases, documents without an associated ACL entry are not ingested. Make sure every document has an ACL defined either through the global ACL file or in its metadata file.

Global ACL configuration file

Create a single centralized file that defines access permissions at the folder level. This provides a streamlined way to manage permissions across large document hierarchies. This method is ideal for organizations with stable permission structures. Any changes to the global file require reindexing the entire affected prefix, which can take hours for knowledge bases with tens of millions of documents. For the file format, see [the section called "Global ACL file structure"](#).

Document-level metadata files

Each document has its own metadata file containing specific access control information. This approach requires you to create and maintain individual metadata files for each document. It enables significantly faster index updates when permissions change because only the affected documents need to be reindexed rather than entire folder structures. For more information about configuring ACLs in metadata files, see [the section called "Document metadata"](#).

Choose the method that best fits your operational needs: centralized management with the global ACL file for simpler administration, or document-level metadata files for faster permission updates and more granular control.

Keep your document-level ACLs current by regularly updating the Amazon S3 ACL configuration to match your organization's access requirements. For more information about common best practices, see [the section called "Best practices for managing ACLs in knowledge bases"](#).

Global ACL file structure

The global file provides centralized access control management at the folder level. Each entry in the file maps an Amazon S3 key prefix to a set of ACL entries that apply to all documents under that prefix.

The global ACL json file uses the following structure:

```
[
```

```

{
  "keyPrefix": "s3://BUCKETNAME/prefix1/",
  "aclEntries": [
    {
      "Name": "user1@example.com",
      "Type": "USER",
      "Access": "ALLOW"
    },
    {
      "Name": "group1",
      "Type": "GROUP",
      "Access": "DENY"
    }
  ]
},
{
  "keyPrefix": "s3://BUCKETNAME/prefix1/document_1.txt",
  "aclEntries": [
    {
      "Name": "user1@example.com",
      "Type": "USER",
      "Access": "ALLOW"
    },
    {
      "Name": "group1",
      "Type": "GROUP",
      "Access": "DENY"
    }
  ]
},
{
  "keyPrefix": "s3://BUCKETNAME/prefix2/",
  "aclEntries": [
    {
      "Name": "user2@example.com",
      "Type": "USER",
      "Access": "ALLOW"
    },
    {
      "Name": "user1@example.com",
      "Type": "USER",
      "Access": "DENY"
    }
  ]
}

```

```
[
  {
    "Name": "group1",
    "Type": "GROUP",
    "Access": "DENY"
  }
]
```

Each entry in the array contains the following fields:

keyPrefix

The Amazon S3 path prefix that the ACL entries apply to. All documents under this prefix inherit the specified permissions.

aclEntries

An array of access control entries, each containing the following fields:

- Name – For USER type, the email address of the user in Quick. For GROUP type, the group name in Quick.
- Type – Either USER or GROUP.
- Access – Either ALLOW or DENY.

Document metadata

You can add metadata to documents in your Amazon S3 bucket to customize chat results and control document-level access. Metadata is additional information about a document, such as its title, creation date, and access permissions.

Amazon Quick supports source attribution with citations. If you specify the `_source_uri` metadata field, the source attribution links in chat results direct users to the configured URL. If you don't specify a `_source_uri`, users can still access source documents through clickable citation links that download the file at query time.

Document metadata location

In Amazon S3, each metadata file can be associated with an indexed document. Your metadata files must be stored in the same Amazon S3 bucket as your indexed files. You can specify a location within the Amazon S3 bucket for your metadata files when configuring your Amazon S3 integration in Amazon Quick.

If you don't specify an Amazon S3 prefix, your metadata files must be stored in the same location as your indexed documents. If you specify an Amazon S3 prefix for your metadata files, they must be in a directory structure parallel to your indexed documents. Amazon Quick looks only in the specified directory for your metadata. If the metadata isn't read, check that the directory location matches the location of your metadata.

The following examples show how the indexed document location maps to the metadata file location. The document's Amazon S3 key is appended to the metadata's Amazon S3 prefix and then suffixed with `.metadata.json` to form the metadata file's Amazon S3 path.

Note

The combined Amazon S3 key, the metadata's Amazon S3 prefix, and the `.metadata.json` suffix must be no more than a total of 1,024 characters. We recommend that your Amazon S3 key is less than 1,000 characters to account for additional characters when combining your key with the prefix and suffix.

Example 1: No metadata path specified

```
Bucket name:
    s3://bucketName
Document path:
    documents
Metadata path:
    none
File mapping
    s3://bucketName/documents/file.txt ->
    s3://bucketName/documents/file.txt.metadata.json
```

Example 2: Metadata path specified

```
Bucket name:
    s3://bucketName
Document path:
    documents/legal
Metadata path:
```

```

    metadata
File mapping
    s3://bucketName/documents/legal/file.txt ->
        s3://bucketName/metadata/documents/legal/file.txt.metadata.json

```

Document metadata structure

You define your document metadata itself in a JSON file. The file must be a UTF-8 text file without a BOM marker. The file name of the JSON file must be `<document>.<extension>.metadata.json`. In this example, `document` is the name of the document that the metadata applies to and `extension` is the file extension for the document. The document ID must be unique in `<document>.<extension>.metadata.json`.

The content of the JSON file uses the following template:

```

{
  "DocumentId": "document ID",
  "Attributes": {
    "_authors": ["author of the document"],
    "_category": "document category",
    "_created_at": "ISO 8601 encoded string",
    "_last_updated_at": "ISO 8601 encoded string",
    "_source_uri": "document URI",
    "_version": "file version",
    "_view_count": number of times document has been viewed
  },
  "AccessControlList": [
    {
      "Name": "user1@example.com",
      "Type": "GROUP | USER",
      "Access": "ALLOW | DENY"
    }
  ],
  "Title": "document title",
  "ContentType": "PDF | HTML | MS_WORD | PLAIN_TEXT | PPT | RTF | XML | XSLT |
MS_EXCEL | CSV | JSON | MD"
}

```

If you provide a metadata path, make sure that directory structure inside the metadata directory exactly matches the directory structure of data file.

For example, if the data file location is at `s3://bucketName/documents/legal/file.txt`, the metadata file location should be at `s3://bucketName/metadata/documents/legal/file.txt.metadata.json`.

All of the attributes and fields are optional, so it's not necessary to include all attributes. However, you must provide a value for each attribute that you want to include; the value can't be empty.

The `_created_at` and `_last_updated_at` metadata fields are ISO 8601 encoded dates. For example, `2012-03-25T12:30:10+01:00` is the ISO 8601 date-time format for March 25, 2012, at 12:30PM (plus 10 seconds) in the Central European Time time zone.

The `AccessControlList` field is an optional array that defines document-level access control. Each entry in the array contains the following fields:

- **Name** – For **USER** type, the email address of the user in Quick. For **GROUP** type, the group name in Quick.
- **Type** – Either **USER** or **GROUP**.
- **Access** – Either **ALLOW** or **DENY**.

Note

To use the `AccessControlList` field, you must enable document-level ACLs when creating the knowledge base. For more information, see [the section called “Document-level ACLs”](#).

Set up and manage Amazon S3 integration

Set up Amazon S3 integration

After your administrator has completed the setup tasks, follow these steps to create your Amazon S3 knowledge base.

1. In the Amazon Quick console, choose **Knowledge**.
2. Under **Amazon S3**, choose **Add** (the plus + button).
3. On the **Connect S3 bucket** page, choose your data source:
 - To reuse an existing Amazon S3 data source, select it from the dropdown. Then choose **Next** to skip to the knowledge base details step.

- To connect a new Amazon S3 bucket, choose **+ Add account** from the dropdown.
4. If you are connecting a new bucket, fill in the connection details:
 - **Name** – A descriptive name for your Amazon S3 integration.
 - **S3 bucket location** – Choose **Quick Suite instance account** to access Amazon S3 buckets in the same AWS account where Amazon Quick is enabled, or choose **Other AWS account** to access buckets in a different account.
 - **S3 bucket URL** – The Amazon S3 bucket path containing your documents. Your Amazon S3 bucket must be in the same region as your Amazon Quick region.

Choose **Next**. The system validates your configuration. If errors occur, review the error message for specific remediation steps.

 Note

If you receive an access error, contact your administrator to verify that your user has the required permissions for the Amazon S3 bucket.

5. On the **Create knowledge base** page, complete the following:
 - **Name** – Enter a descriptive name for your knowledge base.
 - **Description** – Describe the purpose of this knowledge base (optional).
 - **Content** – Choose **Add all content** to sync everything from the bucket, or choose **Add specific content** to specify S3 prefixes for the folders and files you want to include. Filters are case-sensitive.
6. Choose **Next: Additional settings** to configure ACL and metadata options, or choose **Create** to create the knowledge base with default settings.
7. On the **Additional settings** page, configure ACL management and metadata:

 Important

The decision to enable or disable ACLs must be made during knowledge base creation. You cannot change this option after this step. For more information about ACLs, see [the section called “Document-level ACLs”](#).

- To enable document-level ACLs, select **Control document access with ACLs**. When enabled, the following options appear:
 - **Global ACL file location** – Enter the Amazon S3 path to your global ACL file (e.g. acl.json) if you are using a global ACL configuration file for centralized access control management at the folder level.
 - **Metadata files folder location** – Enter the Amazon S3 path to your metadata folder if you are using document-level metadata files that include ACL entries.
 - If your metadata files use the sidecar method (stored in the same folder as the original documents), you can leave both fields blank.
- You can optionally specify a **Metadata files folder location** even without ACLs enabled.

8. Choose **Create**.

After you choose create, the data sync starts automatically.

Manage knowledge bases

After setting up your Amazon S3 integration, you can create and manage knowledge bases from your Amazon S3 content.

Edit existing knowledge bases

You can modify your existing Amazon S3 knowledge bases:

1. In the Amazon Quick console, choose **Knowledge bases**.
2. Select your Amazon S3 knowledge base from the list.
3. Choose the three-dot icon under **Actions**, then choose **Edit knowledge base**.
4. Update your configuration settings as needed and choose **Save**.

Create additional knowledge bases

You can create multiple knowledge bases from the same Amazon S3 integration:

1. In the Amazon Quick console, choose **Knowledge**.
2. Choose your existing Amazon S3 integration from the list.
3. Choose the three-dot icon under **Actions**, then choose **Create knowledge base**.

4. Configure your knowledge base settings and choose **Create**.

For detailed information about knowledge base configuration options, see [the section called "Common configuration settings"](#).

Note

When you create a knowledge base in Amazon Quick, by default only you can get insights from the knowledge base. For shared content, you can provide access to different users and groups by updating the knowledge base permissions. To control document-level access within a knowledge base, see [the section called "Document-level ACLs"](#).

Troubleshooting Amazon S3 integration issues

If you encounter issues connecting to your Amazon S3 bucket, review the following common causes and solutions.

Documents not appearing in ACL-enabled knowledge base

Issue: Documents are not ingested into an ACL-enabled knowledge base.

Solution: For ACL-enabled knowledge bases, documents without an associated ACL entry are not ingested. Verify that every document has an ACL defined either through the global ACL file or in its metadata file. For more information, see [the section called "Document-level ACLs"](#).

Cross-account access not configured

Issue: Your administrator hasn't granted access to use Amazon S3 buckets from other AWS accounts in Amazon Quick.

Solution: Ask your administrator to grant cross-account Amazon S3 access. For more information, see [the section called "Grant Amazon Quick access to Amazon S3 buckets"](#).

Bucket not in approved list

Issue: The bucket you're trying to access hasn't been authorized by your administrator.

Solution:

- Confirm the bucket name is spelled correctly.

- Verify with your administrator that the bucket is included in the approved list.
- Request your administrator to add the bucket to the authorized buckets list if needed.

Insufficient IAM permissions

Issue: Your IAM role or user lacks the necessary permissions to access the Amazon S3 bucket.

Solution:

- Verify your IAM policy includes the required Amazon S3 permissions:
 - `s3:GetObject`
 - `s3:ListBucket`
 - `s3:GetBucketLocation`
 - `s3:GetObjectVersion`
 - `s3:ListBucketVersions`
- Check your own buckets for any explicit Deny statements that might be blocking access.

 Note

The ARN `arn:aws:iam::account-id:role/service-role/aws-quicksight-service-role-v0` is the default service role used when no custom role has been created. If a custom service role exists, contact your administrator to obtain the custom service role ARN and use it instead of the default.

Cross-region restrictions

Issue: The Amazon S3 bucket is located in a different AWS region than your Amazon Quick account or service.

Solution:

- Verify the bucket region matches your Amazon Quick service region.
- Check bucket region using AWS CLI: `aws s3api get-bucket-location --bucket bucket-name`
- Use a bucket in the same region as your service.

Additional troubleshooting steps

- **Test bucket accessibility** using AWS CLI:

```
aws s3 ls s3://bucket-name --profile your-profile
```

- **Review CloudTrail logs** for AccessDenied errors to identify the specific permission issue.
- **Check Amazon S3 Block Public Access settings** - while these typically don't affect authenticated access, verify they're not interfering with your specific use case.
- **Verify bucket ownership** - ensure the bucket exists and you have the correct bucket name.

Limitations

When using Amazon S3 integrations in Amazon Quick, be aware of the following limitations:

- The Amazon S3 bucket must be in the same AWS Region as your Amazon Quick application.
- Each document can have a maximum of 2,500 individual user or group ACL entries.
- Global ACL configuration file maximum size: 100 MB.
- Document metadata file maximum size: 5 MB.

For more information about document-level ACL limitations, see [the section called “Limitations”](#).

Asana integration

Connect Amazon Quick to your Asana workspace to manage projects, tasks, and team collaboration. The Asana connector uses the Asana V2 Model Context Protocol (MCP) server. You can create, update, and manage Asana content without leaving your Amazon Quick environment. Every action runs as the authenticated user, and the connector limits access to what that user can do in Asana. For Amazon Quick subscription requirements, see [the section called “Set up integrations in the console”](#).

Supported authentication methods

The Asana connector supports a single authentication method: **Custom OAuth app**. This method uses user-based Open Authorization (OAuth) to connect to the Asana V2 MCP server.

Note the following behavior when you authenticate with the Asana MCP server:

- The Asana MCP server does not support Dynamic Client Registration. You must pre-register an OAuth app in the Asana developer console before you set up the connector.
- MCP apps do not use scopes. Remove the scope parameter entirely from your authorization URL.
- Tokens issued for MCP apps work only with the MCP server, not with the standard Asana API.
- Access tokens expire after 1 hour. The connector uses refresh tokens to maintain access.

Before you begin

Before you set up the Asana connector, make sure you have the following:

- An Asana account with access to the Asana developer console.
- Permission to create and distribute apps in the Asana workspaces you want to connect.
- A Amazon Quick subscription that meets the integration requirements. For more information, see [the section called “Set up integrations in the console”](#).

Configuring the Asana OAuth MCP app

Create an OAuth MCP app in the Asana developer console, configure its OAuth settings, and configure workspace access before you set up the connector in Amazon Quick.

Create an OAuth MCP app

1. Go to the Asana developer console and sign in.
2. Choose **Create new app**.
3. Enter the app name.
4. For app type, choose **MCP app**.
5. Choose **Create app**.
6. Record the **Client ID** and **Client Secret**. You enter these values when you set up the connector in Amazon Quick.

Configure OAuth settings

1. In the sidebar, choose **OAuth**.
2. Add the Amazon Quick callback URL as the redirect URL. You can copy this value from the connector setup page in Amazon Quick.

Configure workspace access

1. In the sidebar, choose **Manage distribution**.
2. Choose how the app can be used:

Specific workspaces

Choose the individual workspaces where the app can be used.

Any workspace

Allow the app to be used in any workspace.

Important

If you choose **Specific workspaces** but do not choose any workspaces, users receive an error when they try to authenticate.

Setting up the connector in Amazon Quick

After you configure your Asana OAuth MCP app, set up the connector in Amazon Quick.

1. In the Amazon Quick console, choose **Connectors**.
2. Choose the **Create for your team** tab.
3. Find and choose **Asana**.
4. Fill in the following details:
 - **Name** - Enter a descriptive name for your Asana integration.
 - **Description** - Describe the purpose of this integration.
5. Choose the connection type and configure the network type settings.
6. Configure the authentication settings. Use the values from your Asana OAuth MCP app.

Auth Type

Choose **Custom OAuth app**.

Client ID

Enter the client ID from the Asana developer console.

Client Secret

Enter the client secret from the Asana developer console.

Authorization URL

Enter `https://app.asana.com/-/oauth_authorize`. Do not include a scope parameter.

Token URL

Enter `https://app.asana.com/-/oauth_token`.

Redirect URL

This field is pre-filled with the Amazon Quick callback URL. Use this value as the redirect URL in your Asana OAuth app.

7. Choose **Create and continue**.
8. Add users to share the integration with.
9. Choose **Next**.

Managing and troubleshooting

You can edit, share, and delete your Asana integrations after you create them. For more information, see [the section called "Managing existing integrations"](#).

If you run into problems with your Asana integration, use the following guidance to resolve common issues.

This app is not available to your workspace

The user's workspace is not included in the app's distribution settings. In the Asana developer console, choose **Manage distribution** and add the workspace, or choose **Any workspace**.

Invalid scope(s) requested

Remove the scope parameter from the authorization URL. MCP apps do not use scopes.

Authentication failures

Verify that the redirect URL matches the Amazon Quick callback URL, that the client ID and client secret are correct, and that the app is properly configured in **Manage distribution**.

Connection issues

Verify that the access token is included in your requests and has not expired. Access tokens have a 1-hour lifetime. If the token has expired, the connector uses refresh tokens to obtain a new one automatically.

Atlassian Confluence Cloud integration

Use the Atlassian Confluence Cloud integration to perform actions on Confluence content, or to create knowledge bases from Confluence spaces, pages, and blog posts for AI-powered search and Q&A.

Integration capabilities

Connector

Create and update Confluence pages, and search Confluence content through Atlassian API calls.

Knowledge base

Index Confluence pages, blog posts, and space content. Amazon Quick agents can then search and answer questions about the content. Optionally enable document-level access controls to ensure users only see answers from content they are authorized to access in Confluence.

Atlassian Confluence Cloud action integration

Use the Atlassian Confluence Cloud connector to create, update, and manage Confluence pages and spaces directly in Amazon Quick through natural language.

To set up this integration, you first create an OAuth 2.0 (3LO) app in the Atlassian Developer Console and configure its permissions. Then, you create the integration in Amazon Quick and connect it to your Atlassian app. For information about the authentication methods that Amazon Quick supports, see [the section called "Authentication methods"](#).

Before you begin

Make sure you have the following before you set up the integration.

- Atlassian Confluence Cloud.
- Access to the [Atlassian Developer Console](#) to create or manage an OAuth app.

- For subscription requirements, see [the section called “Set up integrations in the console”](#).

Configure the Atlassian Developer Console

If you plan to use user authentication (3LO), create an OAuth 2.0 app in the Atlassian Developer Console before you configure Amazon Quick. Complete all of the following steps before moving to the Amazon Quick console.

If you plan to use service authentication (API Key) only, you can skip this section and proceed to [the section called “Prepare authentication”](#).

For more information about OAuth 2.0 (3LO) apps, see [OAuth 2.0 \(3LO\) apps](#) in the Atlassian developer documentation.

Create an OAuth 2.0 (3LO) app

Amazon Quick uses an Atlassian OAuth 2.0 (3LO) app to authenticate with your Atlassian Cloud product on behalf of your users. Create this app in the Atlassian Developer Console before you configure Amazon Quick.

1. Open the [Atlassian Developer Console](#) and sign in with your Atlassian account.
2. Choose **Create**, then choose **OAuth 2.0 integration**.
3. For **Name**, enter a descriptive name for your integration, for example *your-app-name* connector.
4. Review and accept the Atlassian developer terms.
5. Choose **Create**.

Configure permissions

After you create the OAuth 2.0 app, add the API permissions that Amazon Quick needs to interact with your Atlassian product.

1. From your app in the Atlassian Developer Console, choose **Permissions** in the left navigation.
2. Find the API for your Atlassian product (for example, **Jira API** or **Confluence API**) and choose **Add**. The button changes to **Configure** after the API is added.
3. Choose **Configure**. The scopes page opens with **Classic scopes** and **Granular scopes** tabs.
4. On the **Classic scopes** tab, choose **Edit Scopes**. Select the required classic scopes and choose **Save**.

5. Choose the **Granular scopes** tab, then choose **Edit Scopes**. Select the required granular scopes and choose **Save**.

For the specific scopes required for your integration, see the scopes section that follows.

Required scopes for Confluence

Add the following scopes to your OAuth 2.0 app for the Confluence Cloud action integration.

Classic scope

On the **Classic scopes** tab, choose **Edit Scopes** and select the following scope.

Confluence action integration – classic scope

Scope	Description
search:confluence	Search Confluence content and space summaries.

Granular scopes

Choose the **Granular scopes** tab, then choose **Edit Scopes**. Select the following scopes.

Confluence action integration – granular scopes

Scope	Description
read:page:confluence	View page content.
write:page:confluence	Create and update pages.
read:space:confluence	Access space information.

Configure authorization

Set the callback URL so that Atlassian can redirect users back to Amazon Quick after they authorize the app.

1. From your app in the Atlassian Developer Console, choose **Authorization** in the left navigation.

2. Next to **OAuth 2.0 (3LO)**, choose **Add**.
3. For **Callback URLs**, enter `https://region.quicksight.aws.amazon.com/sn/oauthcallback`. Replace *region* with the AWS Region where your Amazon Quick instance is deployed, for example `us-east-1`.
4. Choose **Save changes**.

Record your credentials

Before you leave the Atlassian Developer Console, confirm that you have the following values. You need them for the Amazon Quick configuration.

1. From your app in the Atlassian Developer Console, choose **Settings** in the left navigation.
2. Under **Authentication details**, copy the **Client ID** and **Secret** values.

Required credentials from Atlassian Developer Console

Value	Where to find it
Client ID	Settings page, under Authentication details
Secret	Settings page, under Authentication details

Prepare authentication

Confluence Cloud connectors support two authentication methods. Gather the required credentials before configuring Amazon Quick.

User authentication (3LO)

If you completed the Atlassian Developer Console setup in the previous section, you should have the following values ready. Enter these when you configure the integration in Amazon Quick.

- **Base URL** – Your Confluence instance URL for API calls. This is not the same URL that users log into. It resembles the following: `https://api.atlassian.com/ex/confluence/yourInstanceId`. To find your instance ID, navigate to `https://your-domain.atlassian.net/_edge/tenant_info`.
- **Client ID** – From the Settings page of your Atlassian OAuth app.

- **Client secret** – From the Settings page of your Atlassian OAuth app.
- **Token URL** – <https://auth.atlassian.com/oauth/token>
- **Authorization URL** – <https://auth.atlassian.com/authorize>
- **Redirect URL** – [https://*region*.quicksight.aws.amazon.com/sn/oauthcallback](https://region.quicksight.aws.amazon.com/sn/oauthcallback)

Service authentication (API Key)

Gather the following information from your Confluence Cloud administrator:

- **API Key** – Confluence API token.
- **Base URL** – Your Confluence instance URL for API calls.
- **Email** – Associated user account email.

Set up the integration in Amazon Quick

After you prepare your authentication credentials, create the integration in Amazon Quick.

1. In the Amazon Quick console, choose **Connectors**.
2. Choose the **Create for your team** tab.
3. Find and choose **Atlassian Confluence Cloud**.
4. On the **Integration type** page, select **Perform actions in Atlassian Confluence Cloud** and choose **Next**.
5. In the **Create integration** form, fill in the following fields:
 - **Name** – Descriptive name for your Confluence integration.
 - **Description** (Optional) – Notes about how this connection will be used.
 - **Connection type** – Choose **Public network**.
6. Under **Authentication settings**, choose your authentication method and fill in the required fields:
 - a. For **User authentication**, configure the following fields:
 - **Base URL** – Your Confluence instance URL for API calls, in the format [https://api.atlassian.com/ex/confluence/*yourInstanceId*](https://api.atlassian.com/ex/confluence/yourInstanceId). To find your instance ID, navigate to [https://*your-domain*.atlassian.net/_edge/tenant_info](https://your-domain.atlassian.net/_edge/tenant_info).
 - **Client ID** – Client ID from the Settings page of your Atlassian OAuth app.
 - **Client secret** – Secret from the Settings page of your Atlassian OAuth app.

- **Token URL** – `https://auth.atlassian.com/oauth/token`
 - **Authorization URL** – `https://auth.atlassian.com/authorize`
 - **Redirect URL** – This field is pre-populated with your Amazon Quick callback URL.
- b. For **Service authentication**, configure the following fields:
- **API Key** – Confluence API token.
 - **Base URL** – Your Confluence instance URL for API calls.
 - **Email** – Associated user account email.
7. Choose **Create and continue**.
8. (Optional) On the **Share integration** page, choose users to share the integration with.

Available actions

After you set up the integration, the following actions are available.

Confluence Cloud available actions

Action	Description
Create Page	Create a new page.
Get Pages	View all pages.
Search	Search content using CQL.
Update Page	Update page content.

Manage and troubleshoot

To edit, share, or delete your integration, see [the section called “Managing existing integrations”](#).

Authentication issues

- **Incorrect app configuration** – Verify the OAuth app in the Atlassian Developer Console includes the required scopes and that the redirect URI matches your Amazon Quick configuration.
- **Expired API token** – If using service authentication, check that the API token has not expired and generate a new one if needed.

- **Incorrect Base URL** – The Base URL for API calls is not the same as the Confluence Cloud login URL. Verify you are using the API URL format: `https://api.atlassian.com/ex/confluence/yourInstanceId`. To find your instance ID, navigate to `https://your-domain.atlassian.net/_edge/tenant_info`.

Common error messages

- **Access denied. You do not have permission to perform this action** – The authenticated user does not have the required permissions in Confluence Cloud. Contact your Confluence Cloud administrator to verify and grant appropriate permissions.
- **OAuth 2.0 authorization failed** – Verify the client ID, client secret, and OAuth scopes are configured correctly in both the Atlassian Developer Console and Amazon Quick.

Atlassian Confluence Cloud knowledge base integration

Use the Atlassian Confluence Cloud knowledge base integration to index Confluence content so that Amazon Quick agents can search and answer questions about it. Amazon Quick supports optional document-level access controls (ACLs) for Confluence knowledge bases. When enabled, users only see answers from documents they are authorized to access in Confluence. For more information, see [the section called “Document-level access controls”](#).

Before you begin

Make sure you have the following before you set up the integration.

- An Atlassian Confluence Cloud instance.
- For Quick subscription requirements, see [the section called “Set up integrations in the console”](#).

If you plan to enable document-level access controls, you must provide Atlassian admin credentials during setup. These credentials allow Amazon Quick to access all user and group information regardless of individual email visibility settings. Without admin credentials, the ACL option is not available during knowledge base creation. To provide admin credentials, you need the following from your Atlassian organization administrator:

- **Atlassian admin API key** – An alphanumeric string generated in the Atlassian admin portal.
- **Organization ID (UUID)** – The unique identifier for your Atlassian organization. You can find this in your browser URL when signed in to the admin portal: `admin.atlassian.com/o/orgId`.

- **Directory ID (UUID)** – The unique identifier for your user directory. Retrieve this value using the Atlassian Admin Workspace API.

For detailed steps to obtain these credentials, see [the section called “Obtain Atlassian admin credentials”](#).

Set up the knowledge base integration

The setup wizard has three steps: authentication, knowledge base creation, and additional settings.

Authenticate your account

1. In the Amazon Quick console, choose **Knowledge**.
2. Under **Knowledge**, choose **Atlassian Confluence Cloud** and choose the Add (plus "+") button.
3. In the **Create Confluence knowledge base** dialog, under **Authentication method**, choose an existing connected account from the **Name** dropdown, or choose **Add account** to create a new connection.
4. If you are adding a new account, complete the following fields:
 - **Name** – A descriptive name for your connection.
 - **Description** (Optional) – Notes about the connection.
 - **Confluence URL** – The URL of your Atlassian site (for example, `https://your-site.atlassian.net`).
5. (Optional) Under **ACL Management**, select **Use Atlassian admin credentials** to provide Atlassian admin credentials. When selected, enter the following:
 - **API key** – The Atlassian admin API key.
 - **OrganizationID (UUID)** – Your Atlassian organization ID.
 - **Directory ID (UUID)** – Your Atlassian user directory ID.

Important

Atlassian admin credentials are required to enable document-level access controls (ACLs) in **Additional settings**. Without admin credentials, the ACL option is disabled. Admin credentials cannot be added after the knowledge base is created. For steps to obtain these credentials, see [the section called “Obtain Atlassian admin credentials”](#).

6. Choose **Next**.
7. Complete the Confluence Cloud authentication popup that appears. Review the requested permissions and choose **Accept** to authorize Amazon Quick to access your Confluence content.

Create the knowledge base

1. Under **Create knowledge base**, enter a **Name** for your knowledge base.
2. (Optional) Enter a **Description** with notes about how the knowledge base will be used.
3. Under **Confluence URLs**, paste the URLs of the Confluence spaces, blogs, or pages that you want to include. Choose **Add** after each URL.

Note

URLs that follow the structure `https://company.atlassian.net/wiki/spaces/space-key/overview` are treated as page URLs, not space URLs. To index an entire space, use the space URL without `/overview` (for example, `https://company.atlassian.net/wiki/spaces/space-key`).

4. Choose **Next: Additional settings** to configure access controls and indexing options, or choose **Create** to create the knowledge base with default settings.

Configure additional settings

1. To enable document-level access controls, select **Control document access with ACLs**. This option is only available if you provided Atlassian admin credentials in **Authentication method**. When enabled, Amazon Quick syncs access control lists from Confluence and verifies each user's permissions at query time. For more information about how access controls work, see [the section called "Document-level access controls"](#).

If you do not enable ACLs, access is controlled at the knowledge base level. Anyone who has access to the knowledge base can get insights from all of the content within it. When ACLs are enabled, access is controlled at the document level. Users who have access to the knowledge base only see answers from documents that they are authorized to access in Confluence.

 Important

ACL management cannot be changed after the knowledge base is created. If you need to change this setting, you must create a new knowledge base.

 Note

If you did not provide Atlassian admin credentials in **Authentication method**, this option is disabled. To enable ACLs, choose **Back** or cancel the wizard and start again with a connected account that includes admin credentials.

2. Under **Multi-media content, file size, and file patterns**, configure the following indexing options as needed:
 - **Visual content in documents** – Index visual content such as images and charts embedded in documents.
 - **Audio files** – Index audio file attachments.
 - **Video files** – Index video file attachments.
 - **Maximum single file size** – The maximum file size for individual attachments (default: 50 MB).
3. Choose **Create**.

Supported content types

- Confluence pages and blog posts
- Page and blog post attachments

When you add a Confluence space URL, Amazon Quick crawls all pages and blog posts within that space. For supported attachment file types and size limits, see [the section called “File size and content limits”](#).

Manage knowledge bases

Edit existing knowledge bases

1. In the Amazon Quick console, choose **Knowledge bases**.
2. Select your Confluence Cloud knowledge base from the list.
3. Choose the three-dot icon under **Actions**, then choose **Edit knowledge base**.
4. Update your configuration settings as needed and choose **Save**.

Troubleshoot

To edit, share, or delete your integration, see [the section called “Managing existing integrations”](#).

For general knowledge base troubleshooting, including sync issues and missing documents, see [the section called “Troubleshooting knowledge bases”](#).

Blocked OAuth app authorization

Symptoms:

- Error message: "Your site admin must authorize this app for the site *instance-name*.atlassian.net before the app can access your account."
- Choosing **Accept** in the consent dialog has no effect.

Cause:

Your Atlassian site administrator has blocked user-installed OAuth apps. When this setting is enabled, only a site or organization administrator can authorize new third-party apps.

Resolution steps:

Use one of the following options to resolve this issue.

- **Option 1 (Recommended): Admin authorizes the app directly**

1. An Atlassian site administrator navigates to Amazon Quick and starts a new knowledge base setup with Confluence Cloud.
2. Because the administrator has site-level permissions, a clean consent screen appears without the error.
3. The administrator chooses **Accept** to install the app.

After the administrator authorizes the app, all other users on the site can connect without issues.

- **Option 2: Temporarily allow user-installed apps** – An administrator goes to `admin.atlassian.com`, navigates to **Apps, Atlassian Apps**, then chooses the link for third-party and Marketplace apps. Under **Settings**, find **User Installed Apps** and toggle to allow user apps. After the user authorizes Amazon Quick, toggle the setting back to block user apps.

Important

Admin authorization applies per Atlassian site, not per organization. If your company has multiple sites (for example, `team-a.atlassian.net` and `team-b.atlassian.net`), each site requires separate authorization.

Note

While user-installed apps are unblocked (Option 2), any user on the site can authorize any OAuth app. Re-enable the block promptly after the user has connected.

Authentication popup fails

Symptoms:

- Authentication popup does not appear or closes immediately.
- Popup appears but fails to complete the OAuth flow.

Resolution steps:

1. Verify that your browser allows popups from the Amazon Quick console domain.
2. Verify that your Confluence Cloud instance is accessible from your network.
3. Try using a different browser or clearing your browser cache.

Missing content in knowledge base

Symptoms:

- Knowledge base sync completes but expected content is not indexed.
- Search results do not include content from specific spaces or pages.
- Only one document is indexed for an entire space.

Resolution steps:

1. Verify that the Confluence Cloud user who authenticated has access to the spaces and pages you selected during setup.
2. Check that the selected content types are supported (pages, blog posts, and attachments).
3. Review the content selection in your knowledge base configuration to confirm the correct spaces and pages are included.
4. Check your Confluence URLs for the /overview suffix. URLs that end with /wiki/spaces/*space-key*/overview are treated as a single page URL, not a full space. If you intended to index the entire space, use the space URL without /overview (for example, <https://company.atlassian.net/wiki/spaces/space-key>).

ACL-related issues

No results from ACL-enabled knowledge base

- Verify that the end user has signed in to Confluence when prompted during their first query. The sign-in is required for real-time permission verification.
- Verify that the end user has access to the relevant Confluence content in their Confluence instance.

Documents skipped during sync

- If sync reports show items with status SKIPPED, verify that the authenticating user has access to the Confluence spaces and pages included in the knowledge base.
- For more information about verifying document access, see [the section called “Check document access \(ACL verification\)”](#).

Document-level access controls

Confluence Cloud knowledge bases optionally support document-level access control. When enabled, Amazon Quick syncs access control lists (ACLs) from Confluence during each crawl and

verifies each user's permissions at query time. Users only see answers from documents they are authorized to access in Confluence.

How it works

When a user queries an Amazon Quick agent that uses a Confluence knowledge base with ACL management enabled, the system enforces access controls in two stages:

1. **Pre-retrieval filtering** – Amazon Quick performs a semantic search against the vector index to find the most relevant document passages. The system applies access control lists that were synced from Confluence during the last crawl. This produces a preliminary set of candidate documents.
2. **Real-time verification** – The system verifies the candidate documents in real time by checking the querying user's current access in Confluence. Only documents the user is currently authorized to access appear in the response.

This two-stage approach provides document-level access control that stays current even when Confluence permissions change between syncs.

Amazon Quick crawls the following Confluence ACL resources:

- **Spaces** – Space permissions apply to all documents in the space by default.
- **Pages** – Pages can be restricted to specific users and groups. Nested pages inherit restrictions from the parent page and can have their own restrictions.
- **Blogs** – Blog posts can be restricted to specific users and groups in the space.
- **Attachments** – Files attached to pages or blog posts inherit the access controls of their parent document.

Enable ACL management

You configure ACL management during knowledge base creation in the **Additional settings** step. Select the **Control document access with ACLs** checkbox to enable it. This option requires Atlassian admin credentials, which you provide during the **Authentication method** step. Without admin credentials, the ACL option is disabled.

⚠ Important

You cannot change ACL management after you create the knowledge base. If you need to change this setting, you must create a new knowledge base.

Atlassian admin credentials allow Amazon Quick to access all user and group information from your Confluence instance, regardless of individual email visibility settings. For steps to obtain admin credentials from your Atlassian organization, see [the section called “Obtain Atlassian admin credentials”](#). For the knowledge base setup steps, see [the section called “Authenticate your account”](#).

For more information about ACL best practices, see [the section called “Best practices for managing ACLs in knowledge bases”](#).

Real-time access verification

When a user submits a query that involves content from an ACL-enabled Confluence knowledge base, Amazon Quick verifies the user's access in real time:

1. The user asks a question in the Quick chat assistant.
2. If the answer involves Confluence content from an ACL-enabled knowledge base, Quick prompts the user to sign in to Confluence.
3. The user signs in with their Confluence email and password and accepts the authorization dialog.
4. Quick uses the user's credentials to verify access to each candidate document in real time against the Confluence API.
5. Only documents the user currently has access to in Confluence appear in the response.

The sign-in is a one-time step. After signing in, users are not prompted again until their session expires.

ℹ Note

If real-time permission verification fails (for example, if the Confluence API is temporarily unavailable), users see an error message prompting them to retry. Amazon Quick does not fall back to cached permissions — real-time verification is mandatory to ensure access control accuracy.

Obtain Atlassian admin credentials

To provide Atlassian admin credentials during Confluence knowledge base setup, your Atlassian organization administrator must complete the following steps.

Get your API key and Organization ID

1. Sign in to the [Atlassian admin portal](#) with administrator permissions.
2. Open the Administration app for your organization. The URL should look like: `https://admin.atlassian.com/o/ORGANIZATION-UUID/overview`.
3. Copy the **Organization ID** from the URL.
4. Choose **Settings**, then choose **API Keys**.
5. Choose **Create API key**.
6. Select the following scopes for the API key:
 - `read:directories:admin`
 - `read:workspaces:admin`
7. Copy and save both the **Organization ID** and **API Key**.

Note

API keys expire. Monitor the expiration date and update your data source credentials before the key expires.

Get your Directory ID

Use the Atlassian Admin Workspace API to retrieve your Directory ID.

1. Run the following command, replacing **ORGANIZATION-ID** with your Organization ID and **API-KEY** with your API key:

```
curl --request POST \  
  --url 'https://api.atlassian.com/admin/v2/orgs/ORGANIZATION-ID/workspaces' \  
  --header 'Authorization: Bearer API-KEY' \  
  --header 'Accept: application/json' \  
  --header 'Content-Type: application/json' \  
  --data '{
```

```
"query": {
  "field": {
    "name": "attributes.type",
    "values": ["confluence"]
  }
},
"limit": 20
}'
```

This query filters results to Confluence workspaces only, so you don't need to page through all workspace types.

2. In the API response, find the workspace entry that matches your Confluence Cloud instance. Verify the workspace name matches your instance.
3. Copy the `directory` value from the `attributes` section. This is your **Directory ID**.

For more information about Atlassian admin API scopes, see [Atlassian API scopes documentation](#). For API details, see [Atlassian Admin Workspace API reference](#).

Limitations

- **Atlassian admin credentials required** – Document-level access controls require Atlassian admin credentials, which must be provided during knowledge base setup. You cannot enable ACLs without admin credentials, and you cannot add admin credentials after the knowledge base is created.
- **Real-time verification is mandatory** – Amazon Quick always performs real-time permission checks at query time for ACL-enabled knowledge bases. If the Confluence API is unavailable, queries that involve ACL-protected content return an error rather than falling back to cached permissions.

For general ACL limitations and best practices, including ACL permanence, research compatibility, and email address management, see [the section called “Best practices for managing ACLs in knowledge bases”](#).

Next steps

To verify document-level access controls and troubleshoot permission issues, see [the section called “Check document access \(ACL verification\)”](#). For information about setting up the Confluence knowledge base integration, see [the section called “Set up the knowledge base integration”](#).

Atlassian Jira Cloud integration

Use the Atlassian Jira Cloud connector to create, update, search, and manage Jira issues, projects, sprints, and users directly in Amazon Quick through natural language.

Setting up this integration involves two steps. First, you create an OAuth 2.0 (3LO) app in the Atlassian Developer Console and configure its permissions. Then, you create the integration in Amazon Quick and connect it to your Atlassian app. For information about the authentication methods that Amazon Quick supports, see [the section called "Authentication methods"](#).

Before you begin

Make sure you have the following before you set up the integration.

- Atlassian Jira Cloud.
- Access to the [Atlassian Developer Console](#) to create or manage an OAuth app.
- For subscription requirements, see [the section called "Set up integrations in the console"](#).

Configure the Atlassian Developer Console

If you plan to use user authentication (3LO), create an OAuth 2.0 app in the Atlassian Developer Console before you configure Amazon Quick. Complete all of the following steps before moving to the Amazon Quick console.

If you plan to use service authentication (API Key) only, you can skip this section and proceed to [the section called "Set up the integration in Amazon Quick"](#).

For more information about OAuth 2.0 (3LO) apps, see [OAuth 2.0 \(3LO\) apps](#) in the Atlassian developer documentation.

Create an OAuth 2.0 (3LO) app

Amazon Quick uses an Atlassian OAuth 2.0 (3LO) app to authenticate with your Atlassian Cloud product on behalf of your users. Create this app in the Atlassian Developer Console before you configure Amazon Quick.

1. Open the [Atlassian Developer Console](#) and sign in with your Atlassian account.
2. Choose **Create**, then choose **OAuth 2.0 integration**.

3. For **Name**, enter a descriptive name for your integration, for example *your-app-name* connector.
4. Review and accept the Atlassian developer terms.
5. Choose **Create**.

Configure permissions

After you create the OAuth 2.0 app, add the API permissions that Amazon Quick needs to interact with your Atlassian product.

1. From your app in the Atlassian Developer Console, choose **Permissions** in the left navigation.
2. Find the API for your Atlassian product (for example, **Jira API** or **Confluence API**) and choose **Add**. The button changes to **Configure** after the API is added.
3. Choose **Configure**. The scopes page opens with **Classic scopes** and **Granular scopes** tabs.
4. On the **Classic scopes** tab, choose **Edit Scopes**. Select the required classic scopes and choose **Save**.
5. Choose the **Granular scopes** tab, then choose **Edit Scopes**. Select the required granular scopes and choose **Save**.

For the specific scopes required for your integration, see the scopes section that follows.

Configure API permissions

Add the following scopes to your OAuth 2.0 app for the Jira Cloud action integration.

Classic scopes

On the **Classic scopes** tab, choose **Edit Scopes** and select the following scopes.

Jira action integration – classic scopes

Scope	Description
read:jira-work	Read Jira project and issue data, search for issues, and objects associated with issues like attachments and worklogs.
manage:jira-project	Create and edit project settings and create new project-level objects (for example, versions and components).

Scope	Description
manage:jira-configuration	Take Jira administration actions (for example, create projects and custom fields, view workflows, manage issue link types).
read:jira-user	View user information in Jira that the user has access to, including usernames, email addresses, and avatars.
write:jira-work	Create and edit issues in Jira, post comments as the user, create worklogs, and delete issues.
manage:jira-webhook	Fetch, register, refresh, and delete dynamically declared Jira webhooks.

Granular scopes

Choose the **Granular scopes** tab, then choose **Edit Scopes**. Use the search bar to find the scopes below. For example, search for `sprint:jira-software` to find sprint-related scopes.

Jira action integration – granular scopes

Scope	Description
read:board-scope:jira-software	Read board configurations.
read:sprint:jira-software	Read sprint information.
write:sprint:jira-software	Create and modify sprints.
delete:sprint:jira-software	Delete sprints.
write:board-scope:jira-software	Manage board configurations.
read:project:jira	Read project details.

Configure authorization

Set the callback URL so that Atlassian can redirect users back to Amazon Quick after they authorize the app.

1. From your app in the Atlassian Developer Console, choose **Authorization** in the left navigation.
2. Next to **OAuth 2.0 (3LO)**, choose **Add**.
3. For **Callback URLs**, enter `https://region.quicksight.aws.amazon.com/sn/oauthcallback`. Replace *region* with the AWS Region where your Amazon Quick instance is deployed, for example `us-east-1`.
4. Choose **Save changes**.

Record your credentials

Before you leave the Atlassian Developer Console, confirm that you have the following values. You need them for the Amazon Quick configuration.

1. From your app in the Atlassian Developer Console, choose **Settings** in the left navigation.
2. Under **Authentication details**, copy the **Client ID** and **Secret** values.

Required credentials from Atlassian Developer Console

Value	Where to find it
Client ID	Settings page, under Authentication details
Secret	Settings page, under Authentication details

Set up the integration in Amazon Quick

After you prepare your authentication credentials, create the integration in Amazon Quick.

1. In the Amazon Quick console, choose **Connectors**.
2. Choose the **Create for your team** tab.
3. Find and choose **Atlassian Jira Cloud**.
4. In the **Create integration** wizard, fill in the following fields:

- **Name** – Descriptive name for your Jira integration.
 - **Description** (Optional) – Notes about how this connection will be used.
 - **Connection type** – Choose **Public network**.
5. Under **Authentication settings**, choose your authentication method and fill in the required fields:
- a. For **User authentication**, configure the following fields:
 - **Base URL** – Your Jira instance URL for API calls. This is not the same URL that users log into. It resembles the following: `https://api.atlassian.com/ex/jira/yourInstanceId`. To find your instance ID, navigate to `https://your-domain.atlassian.net/_edge/tenant_info`.
 - **Client ID** – Client ID from the Settings page of your Atlassian OAuth app.
 - **Client secret** – Secret from the Settings page of your Atlassian OAuth app.
 - **Token URL** – `https://auth.atlassian.com/oauth/token`
 - **Authorization URL** – `https://auth.atlassian.com/authorize`
 - **Redirect URL** – This field is pre-populated with your Amazon Quick callback URL.
 - b. For **Service authentication**, configure the following fields:
 - **API Key** – Jira API token.
 - **Base URL** – Your Jira instance URL for API calls.
 - **Email** – Associated user account email.
6. Choose **Create and continue**.
7. (Optional) On the **Share integration** page, choose users to share the integration with.

Important

Jira Cloud may return HTTP 200 success responses even when API tokens are revoked or improperly configured. For more information, see [JRACLOUD-82932](#). If your integration appears to connect successfully but actions fail unexpectedly, verify that your API token is valid and has not been revoked.

Available actions

After you set up the integration, the following actions are available.

Jira Cloud available actions

Action	Description
Add Attachment	Add an attachment to an issue.
Add Comment	Add new comment.
Change Issue Status	Change task status of an issue.
Create Issue	Create new issue or subtask.
Create Project	Create new project.
Create Sprint	Create a sprint in a project.
Delete Comment	Remove comment.
Delete Issue	Delete an issue in a project.
Delete Project	Remove project.
Delete Sprint	Delete a sprint in a project.
Edit Issue	Modify issue.
Find Users	Search for a Jira user.
Get All Labels	View all labels.
Get All Users	List all Jira users.
Get Attachment Content	View the contents of an attachment.
Get Comments	View issue comments.
Get Issue	View details of an issue in a project.
Get Issue Types For Project	View project issue types.

Action	Description
Get Priorities	View issue priorities.
Get Project	View project details.
Get Sprint	View details of a sprint in a project.
Move Issues to Backlog	Move issues to backlog.
Move Issues To Sprint And Rank	Assign an issue to a sprint.
Search Issues	Search for issues.
Search Projects	Find visible projects.
Search Statuses	Search issue statuses.
Update Comment	Edit comment.
Update Project	Modify project.
Update Sprint	Update sprint details.

Note

The actions you can use depend on the permissions configured in your Jira Cloud instance and your authentication method.

Manage and troubleshoot

To edit, share, or delete your integration, see [the section called “Managing existing integrations”](#).

Authentication issues

- **Incorrect app configuration** – Verify the OAuth app in the Atlassian Developer Console includes the required scopes and that the redirect URI matches your Amazon Quick configuration.

- **Expired or revoked API token** – If using service authentication, check that the API token has not expired or been revoked. Due to a known Jira Cloud behavior ([JRACLOUD-82932](#)), the integration may appear to connect successfully even with invalid tokens.
- **Incorrect Base URL** – The Base URL for API calls is not the same as the Jira Cloud login URL. Verify you are using the API URL format: `https://api.atlassian.com/ex/jira/yourInstanceId`. To find your instance ID, navigate to `https://your-domain.atlassian.net/_edge/tenant_info`.

Common error messages

- **Access denied. You do not have permission to perform this action** – The authenticated user does not have the required permissions in Jira Cloud. Contact your Jira Cloud administrator to verify and grant appropriate permissions.
- **OAuth 2.0 authorization failed** – Verify the client ID, client secret, and OAuth scopes are configured correctly in both the Atlassian Developer Console and Amazon Quick.

AWS Agent Registry integration

With Amazon Quick, you can connect to the [AWS Agent Registry](#). Use this integration to surface the remote Model Context Protocol (MCP) servers registered in your Agent Registry as native Quick connectors. Your organization can maintain a single source of truth for its MCP servers in the Agent Registry while giving Quick users governed access to them.

As an account administrator, you link one Agent Registry to your Quick account from the Quick administration console. After the registry is linked, you see the registry's remote MCP records as pre-configured connector cards on the **Connectors** page. You can then create a connector from one. The connector works like any other Quick connector across Quick Chat, agents, apps, Amazon Quick Flows, and Amazon Quick Research.

Everything happens through the Quick administration console and the **Connectors** page. There is no separate command line or SDK to install. Records are read from the Agent Registry when the **Connectors** page loads. Nothing from the Agent Registry is copied into Quick.

What is supported

The following table shows which Agent Registry records and configurations Amazon Quick supports.

AWS Agent Registry support in Amazon Quick

Supported	Not supported
MCP records that use the <code>mcpServer</code> descriptor.	Skill and custom record types.
Agent records accessed through the MCP protocol (records that use the <code>mcpServer</code> descriptor).	Agent-to-agent (A2A) descriptors.
Remote MCP servers whose record specifies a remote server URL at <code>descriptors.mcpServer.data.remotes[].url</code> .	Local MCP endpoints, such as <code>npx</code> , <code>docker</code> , and local or standard input/output (stdio) servers.

For more information about the account, Region, authorization, and status requirements that a registry must meet, see [the section called “Prerequisites”](#).

Prerequisites

Before an administrator can link a registry, make sure that all of the following are in place:

- **Registry location** – An AWS Agent Registry exists in the same AWS account and the same AWS Region as your Quick account. Cross-account and cross-Region registries are not supported.
- **Registry configuration** – The registry uses AWS IAM (AWS_IAM) authorization with Signature Version 4 (SigV4) and is in the READY status. Registries that use JSON Web Token (JWT) authorization are not supported and do not appear in the Quick administration console.
- **Registry namespace** – Only registries in the `agent-registry` namespace are supported. Older Amazon Bedrock AgentCore registries in the `bedrock-agentcore` namespace are not supported.
- **One registry for each account** – You can connect one Agent Registry for each Quick account.
- **Administrator access** – The administrator signs in to the Quick administration console with an IAM identity that has permission to list and view registries. For more information, see [the section called “IAM permissions”](#).

IAM permissions

The integration uses two distinct sets of permissions that apply to different steps. Keep them separate.

Administrator identity (to browse and select a registry)

The IAM identity that the administrator uses to sign in to the administration console needs the following permissions on the registry. Without these permissions, the **AWS Agent Registry** page cannot list your registries.

- `agent-registry:ListRegistries`

Service role (to load records on the Connectors page)

How you satisfy this requirement depends on which service role your account uses:

- **Quick-managed role** (`aws-quicksight-agent-registry-role-v0`) – Quick grants these permissions automatically when the administrator enables the registry. No action is required.
- **Customer-managed service role** – An administrator must manually add the following permissions to the role for the connected registry:
 - `agent-registry:ListDiscoverableRegistryRecords`
 - `agent-registry:SearchDiscoverableRegistryRecords`
 - `agent-registry:GetDiscoverableRegistryRecord`

If a customer-managed role is missing these permissions, the **Connectors** page cannot load registry records. Quick detects the customer-managed role and displays the role ARN and the required permissions so that the administrator knows exactly what to add. For a complete policy example, see [the section called “IAM policy reference”](#).

Link a registry (administrator)

After the prerequisites and IAM permissions are in place, you link the registry as an account administrator.

1. In the Quick administration console, go to **Manage account** -> **Permissions** -> **AWS Agent Registry**.
2. Review the list of AWS Agent Registries in your account. Only registries that meet the requirements in [the section called “Prerequisites”](#) appear.

3. Turn on the toggle for the registry that you want to connect. Only one registry can be active for each account.
4. Confirm the selection. Quick grants the account's service role the read-only Agent Registry permissions that it needs and links the registry to your Quick account.

Note

When your account uses the Quick-managed service role, linking the registry is the entire administrator action. You do not need to create a role or edit a trust policy. Quick sets up the trust relationship and permissions for you when you enable the registry.

Create a connector (user)

After an administrator links a registry, you can create a connector from it.

1. Open the **Connectors** page.
2. Choose the **Create for your team** tab. Registry-sourced connectors appear on this tab, not on the **Available** tab.
3. Find the MCP connector sourced from the registry by searching or browsing. Each connector comes pre-populated with its MCP server URL, name, and description.
4. Choose the connector to begin setup.
5. Complete the **Authenticate** step by configuring or signing in with the required credentials, such as OAuth. The Agent Registry does not store authentication credentials, so each user provides or authenticates with their own credentials.
6. On the **Manage permissions** step, choose which tools to turn on or turn off.
7. Choose **Create and continue**.
8. Share the connector with the appropriate user groups.

Note

Registry-sourced connectors are classified as **Custom MCP** connectors. This classification matters for settings such as [custom permissions](#).

After setup

After a connector is created, it is available across Quick Chat, agents, apps, Amazon Quick Flows, and Amazon Quick Research, like any other Quick connector.

The connector is an independent Quick resource. Removing the registry link in the administration console does not delete connectors that were already created from it.

Manage connectors

To edit, share, or delete a connector that you created from the Agent Registry, see [the section called “Managing existing integrations”](#).

Troubleshooting

- **Registries do not appear on the AWS Agent Registry page** – Confirm that the administrator's IAM identity has the `agent-registry:ListRegistries` permission. Also confirm that the registry uses AWS IAM authorization (registries that use JWT authorization are not listed), is in the READY status, and is in the same AWS account and Region as your Quick account.
- **The Connectors page cannot load registry records** – If your account uses a customer-managed service role, confirm that the role has the required `agent-registry` permissions for the connected registry. Quick displays the role ARN and the required permissions so that you know what to add. For the policy, see [the section called “IAM policy reference”](#).
- **A record is Approved but no connector card appears** – Quick reads the MCP server endpoint only from `descriptors.mcpServer.data.remotes[].url`. Verify that the field is populated:

```
aws agent-registry-control get-registry-record \  
  --registry-id <registry-id> --record-id <record-id> --region <region>
```

If the descriptor has no data, or data contains no `remotes[].url`, the record is not surfaced as a connector. A record can reach Approved status in this state. Populate `descriptors.mcpServer.data` with the MCP server definition, including `remotes[].url`. If you set data through the API rather than the console, you must also supply `dataSchemaVersion`, or the request fails with `ValidationException: Schema validation failed`.

Limitations

- Only remote MCP servers are supported. Local MCP endpoints, such as npx, docker, and local or stdio servers, are not supported.
- Only MCP records and agent records accessed through the MCP protocol are supported. Agent-to-agent (A2A) descriptors, skill records, and custom records are not supported.

For more information about account, Region, authorization, and registry-count requirements, see [the section called “Prerequisites”](#).

IAM policy reference

If your account uses the Quick-managed role (aws-quicksight-agent-registry-role-v0), you do not need this policy. Quick provisions the role and attaches the required Agent Registry permissions automatically when the administrator enables a registry. This policy applies only when your account uses a customer-managed service role.

If your account uses a customer-managed service role, attach a policy with the following permissions, scoped to the registry that you are connecting. Replace the example resource ARN with your registry's ARN.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "agent-registry:ListDiscoverableRegistryRecords",
        "agent-registry:SearchDiscoverableRegistryRecords",
        "agent-registry:GetDiscoverableRegistryRecord"
      ],
      "Resource": [
        "arn:aws:agent-registry:us-west-2:111122223333:registry/AbCdEf012345",
        "arn:aws:agent-registry:*:111122223333:registry/AbCdEf012345/record/*"
      ]
    }
  ]
}
```

BambooHR integration

Connect Amazon Quick to your BambooHR system to manage employee data, time-off requests, and HR processes. You can create, update, and manage HR content without leaving your Amazon Quick environment. For Amazon Quick subscription requirements, see [the section called “Set up integrations in the console”](#).

What you can do

With BambooHR integration, you can perform actions within your BambooHR systems through the BambooHR API.

Connector

Create, update, and manage employee records, time-off requests, and other HR processes through the BambooHR API.

Set up BambooHR integration

Follow these steps to connect Amazon Quick to your BambooHR system.

1. In the Amazon Quick console, choose **Connectors**.
2. Choose the **Create for your team** tab.
3. Find and choose **BambooHR**.
4. Complete the integration details:
 - **Name** - Enter a descriptive name for your BambooHR integration.
 - **Description** (Optional) - Describe the purpose of this integration.
5. Choose connection type (user or service authentication).
6. Complete the connection settings based on authentication method:
 - **For User authentication (OAuth):**
 - **Base URL** - Your BambooHR instance URL.
 - **Client ID** - OAuth client identifier.
 - **Client Secret** - OAuth client secret.
 - **Token URL** - OAuth token endpoint.
 - **Auth URL** - OAuth authorization endpoint.

- **Redirect URL** - OAuth callback URL.

Required OAuth scopes: When you create your BambooHR OAuth application, configure the scopes you need for your use case. Common scopes include:

- `employee` - Access employee information.
- `employee.write` - Write employee information.
- `time_off` - Access time off information.
- `time_off.write` - Write time off information.
- `company:info` - Access company information.
- `payroll` - Access payroll data.

 Note

Additional scopes may be required depending on the specific BambooHR actions you plan to use. Consult your BambooHR administrator for the complete list of available scopes.

- **For Service authentication (API Key):**
 - **Base URL** - Your BambooHR instance URL.
 - **API Key** - Your BambooHR API key.
 - **Email** - Email address associated with the API key.

7. Select **Create and continue**.
8. Select users to share the integration with.
9. Click **Next**.

Manage BambooHR integration

You can perform these management tasks for your BambooHR integration:

Edit integration settings

To modify your BambooHR integration settings:

1. In the Amazon Quick console, choose **Connectors**.
2. Choose **BambooHR** from the connector list.

3. Select your integration from the list and choose **Edit**.
4. Modify your integration settings as needed.
5. Choose **Save changes**.

Share integration

You can share your BambooHR connector with other users in your organization.

1. From the BambooHR integration details page, choose **Share**.
2. Configure your sharing options and permission levels.
3. Choose **Share integration**.

Delete integration

To permanently remove your BambooHR integration:

1. From the BambooHR integration details page, choose **Delete**.
2. Review the deletion impact, including any workflows or automations using this integration.
3. Type the integration name to confirm deletion.
4. Choose **Delete integration**.

What happens next

After you complete the setup, your BambooHR integration appears in the integrations list. You can use it in Amazon Quick workflows, automations, and AI agents to perform HR-related tasks.

Troubleshoot BambooHR integration

If you encounter issues with your BambooHR integration, try these solutions:

Authentication failures

Verify that your BambooHR credentials are correct and that the API key or OAuth application has the necessary permissions.

Connection timeouts

Check that your BambooHR instance URL is correct and accessible from Amazon Quick.

Permission errors

Ensure that the authenticated user or API key has the required permissions to perform the requested HR operations in BambooHR.

Action execution failures

Review the action parameters and ensure they match the expected format for BambooHR API calls.

Box integration

Use the Box integration to perform actions on Box content, or to create knowledge bases from Box files and folders for AI-powered search and Q&A.

Integration capabilities

The Box integration offers two connector types.

Connector

Access Box AI features, manage files and folders, and organize content in hubs directly in Amazon Quick through natural language.

Knowledge base

Index Box files and folders. Amazon Quick agents can then search and answer questions about the content.

Box action integration

With the Box connector, you can access Box AI features, manage files and folders, and organize content in hubs directly in Amazon Quick through natural language.

Amazon Quick supports the following authentication methods for Box Agent.

- **Default OAuth app** – Uses an AWS-managed OAuth application. No additional credentials are needed. You authenticate directly with your Box account.
- **Custom OAuth app** – Uses a customer-managed OAuth 2.0 application registered in the Box Developer Console. This option gives your organization full control over the OAuth configuration.

For more information about the authentication methods that Amazon Quick supports, see [the section called "Authentication methods"](#).

Before you begin

Make sure that you have the following before you set up the integration.

- A Box account with access to files and folders.
- For **Custom OAuth app**: Access to the Box Developer Console to create a custom app. For more information, see [OAuth 2.0 setup](#) in the Box documentation.
- For Amazon Quick subscription requirements, see [the section called "Set up integrations in the console"](#).

Setting up the connector in Amazon Quick

Connect from the Available tab

If you want to use Default OAuth app authentication, you can connect directly from the **Available** tab without additional configuration.

1. In the Amazon Quick console, choose **Connectors**.
2. On the **Available** tab, find **Box Agent** and choose **Connect**.
3. Complete the Box sign-in flow and grant the requested permissions.

To configure a connector with Custom OAuth app instead, use the **Create for your team** tab as described below.

Create from the Create for your team tab

If you use Custom OAuth app authentication, create a custom app in the Box Developer Console before configuring Amazon Quick. For setup instructions, see [OAuth 2.0 setup](#) in the Box documentation.

1. In the Amazon Quick console, choose **Connectors**.
2. Choose the **Create for your team** tab.
3. Find and choose **Box Agent**.

 **Note**

If a Box Agent connector already exists, a dialog appears with your existing connectors. To use an existing connector, choose it. To create a new one, choose **No, create new**.

4. Enter a **Name** for your connector. Optionally, choose **+ Add Description** to add a description.
5. For **Connection type**, choose **Public network**.
6. For **OAuth Configuration**, choose one of the following:

a. For **Default OAuth app**:

No additional credentials are needed. Choose **Next** to continue.

b. For **Custom OAuth app**, configure the following fields:
 - **Base URL** (Optional) – The Box API base URL. Example: `https://mcp.box.com`
 - **Client ID** – The client ID from your Box app.
 - **Client secret** – The client secret from your Box app.
7. Choose **Next**.
8. A Box authorization window opens. Review the requested permissions and choose **Grant access to Box**.
9. On the **Review** page, review the available actions for the connector. Choose **Next**.
10. On the **Publish** page, choose who can access the connector. You can enable access for everyone in your organization or search for specific teams or groups.
11. Choose **Publish**.

Available actions

After you set up the connector, the following actions are available.

Box Agent available actions

Category	Action	Description
AI	AI Q&A (Hub)	Ask a question to a Box Hub using Box AI.
AI	AI Q&A (Single File)	Ask a question to a single file using Box AI.

Category	Action	Description
AI	AI Q&A (Multiple Files)	Ask a question to multiple files using Box AI.
AI	AI Extract (Freeform)	Extract metadata from files using Box AI in freeform format.
AI	AI Extract (Structured)	Extract structured metadata from files using Box AI based on custom fields or a metadata template.
Files	Get File Content	Extract and return file content in text format for files under 50 MB.
Files	Get File Details	Get comprehensive file information including metadata, permissions, and version details.
Files	Update File Properties	Update file metadata including name, description, tags, and collections.
Files	Upload File	Upload a new file to Box.
Files	Upload File Version	Upload a new version of an existing file.
Files	Search Files	Search for files using keywords with support for metadata filters and file extension filtering.
Folders	Create Folder	Create a new folder in Box.
Folders	Get Folder Details	Retrieve comprehensive folder information including metadata and collaboration settings.
Folders	List Folder Content	List files, folders, and web links in a folder.
Folders	Update Folder Properties	Update folder metadata including name, description, tags, and collections.
Folders	Search Folders	Search for folders by name using keyword matching.

Category	Action	Description
Hubs	Create Hub	Create a new hub for organizing content.
Hubs	Get Hub Details	Retrieve information about a specific hub.
Hubs	Get Hub Items	Retrieve items within a hub including files, folders, and web links.
Hubs	List Hubs	List all hubs accessible to the current user.
Hubs	Update Hub	Update the title or description of a hub.
Hubs	Add Items to Hub	Add files, folders, or web links to a hub.
Comments and tasks	Create File Comment	Create a new comment on a file.
Comments and tasks	List File Comments	List all comments on a file.
Comments and tasks	List Tasks	List all tasks associated with a file.
Users	Who Am I	Return information about the currently authenticated user.

Manage and troubleshoot

For instructions on editing, sharing, or deleting your connector, see [the section called “Managing existing integrations”](#).

- **Authentication popup fails** – Verify that your browser allows popups from the Amazon Quick console domain.
- **Invalid client credentials (Custom OAuth app)** – Verify that the Client ID and Client secret match the values in your Box Developer Console app.
- **Insufficient permissions** – Verify that your Box account has access to the files, folders, or hubs you are trying to access.

Box knowledge base integration

Use the Box knowledge base integration to index Box content so that Amazon Quick agents can search and answer questions about it. You can sync all of the Box content that you have access to, or choose specific Box folders and files. You can also configure advanced indexing settings when needed.

Access control limitation

This user-managed setup does not support document-level access controls (ACLs). ACLs are mechanisms that control which users can access specific documents.

Before you begin

Make sure you have the following before you set up the integration.

- A Box account with access to the files and folders that you want to include in the knowledge base.
- For Quick subscription requirements, see [the section called “Set up integrations in the console”](#).

Setting up the knowledge base integration

The setup wizard has two steps: create the knowledge base, and configure optional additional settings.

Creating the knowledge base

1. Sign in to the Amazon Quick console, and then choose **Knowledge**.
2. Under **Knowledge bases**, find **Box**, and then choose the **Add (+)** icon.
3. In the **Create Box knowledge base** dialog, under **Connected account**, choose **Sign in to Box**. Complete the Box sign-in flow and grant the requested permissions.

Note

The Box knowledge base integration uses Default OAuth app authentication. You authenticate directly with your Box account, and no additional credentials are needed.

4. Under **Create knowledge base**, enter a **Name** for your knowledge base.
5. (Optional) Enter a **Description** with notes about how the knowledge base will be used.
6. Under **Content**, choose which Box content to include:
 - **Add all content** – Includes all of the Box content that you have access to.
 - **Add specific content** – Includes only the Box folders and files that you specify. In the **Box file or folder URLs** field, enter a Box file or folder URL (for example, `https://app.box.com/file/987654321`) and choose **Add**. Repeat for each URL that you want to include.
7. Choose **Next: Additional settings** to configure indexing options, or choose **Create** to create the knowledge base with default settings.

Configuring additional settings

In the **Additional settings (optional)** step, you can configure indexing options under **Multi-media content, file size, and file patterns**.

Under **Advanced indexing**, configure the following options as needed:

- **Visual content in documents** – Extracts and indexes visual content such as images and charts embedded in documents. This option is enabled by default.
- **Audio files** – Transcribes and indexes audio files.
- **Video files** – Transcribes and indexes video files.
- **Maximum single file size** – The maximum file size for individual files, in MB (default: 500 MB).

Choose **Create** to create the knowledge base. After you choose **Create**, the data sync starts automatically.

Managing knowledge bases

Editing existing knowledge bases

1. Sign in to the Amazon Quick console, and then choose **Knowledge bases**.
2. Select your Box knowledge base from the list.
3. Choose the **More options** icon (:) under **Actions**, then choose **Edit knowledge base**.
4. Update your configuration settings as needed and choose **Save**.

Troubleshoot the Box knowledge base integration

To edit, share, or delete your integration, see [the section called “Managing existing integrations”](#).

For general knowledge base troubleshooting, including sync issues and missing documents, see [the section called “Troubleshooting knowledge bases”](#).

Canva integration

With Canva integration in Amazon Quick, you can create, edit, and manage designs and visual content through MCP server connectivity. This integration provides action capabilities for design operations and content creation.

What you can do

Canva integration provides connector capabilities through MCP server connectivity:

- Create new designs from templates
- Edit existing designs and presentations
- Manage design assets and media
- Export designs in various formats
- Share designs and collaborate with team members

Available tools

The Canva MCP server typically provides these tools:

- `comment-on-design` - Add a comment on a Canva design. You need to provide the design ID and the message text. The comment will be added to the design and visible to all users with access to the design.
- `create-design-from-candidate` - Create a new Canva design from a generation job candidate ID. This converts an AI-generated design candidate into an editable Canva design. If successful, returns a design summary containing a design ID that can be used with the `editing_transaction_tools`. To make changes to the design, first call this tool with the `candidate_id` from generate-design results, then use the returned `design_id` with `start-editing-transaction` and subsequent editing tools.
- `create-folder` - Create a new folder in Canva. You can create it at the root level or inside another folder.

- `export-design` - Export a Canva design, doc, presentation, whiteboard, videos and other Canva content types to various formats (PDF, JPG, PNG, PPTX, GIF, MP4). You should use the `get-export-formats` tool first to check which export formats are supported for the design. This tool provides a download URL for the exported file that you can share with users. Always display this download URL to users so they can access their exported content.
- `generate-design` - Generate designs with AI. Use the 'query' parameter to tell AI what you want to create.
- `get-design` - Get detailed information about a Canva design, such as a doc, presentation, whiteboard, video, or sheet. This includes design owner information, title, URLs for editing and viewing, thumbnail, created/updated time, and page count. This tool doesn't work on folders or images. You must provide the design ID, which you can find by using the `search-designs` or `list-folder-items` tools.
- `get-design-content` - Get the text content of a doc, presentation, whiteboard, social media post, sheet, and other designs in Canva. Use this when you only need to read text content without making changes. IMPORTANT: If the user wants to edit, update, change, translate, or fix content, use `start-editing-transaction` instead as it shows content AND enables editing. You must provide the design ID, which you can find with the `search-designs` tool. When given a URL to a Canva design, you can extract the design ID from the URL. Do not use web search to get the content of a design as the content is not accessible to the public. Example URL: `https://www.canva.com/design/{design_id}`.
- `get-design-pages` - Get a list of pages in a Canva design, such as a presentation. Each page includes its index and thumbnail. This tool doesn't work on designs that don't have pages (e.g. Canva docs). You must provide the design ID, which you can find using tools like `search-designs` or `list-folder-items`. You can use 'offset' and 'limit' to paginate through the pages. Use `get-design` to find out the total number of pages, if needed.
- `get-export-formats` - Get the available export formats for a Canva design. This tool lists the formats (PDF, JPG, PNG, PPTX, GIF, MP4) that are supported for exporting the design. Use this tool before calling `export-design` to ensure the format you want is supported.
- `import-design-from-url` - Import a file from a URL as a new Canva design
- `list-comments` - Get a list of comments for a particular Canva design. Comments are discussions attached to designs that help teams collaborate. Each comment can contain replies, mentions, and can be marked as resolved or unresolved. You need to provide the design ID, which you can find using the `search-designs` tool. Use the continuation token to get the next page of results, if needed. You can filter comments by their resolution status (resolved or unresolved) using the `comment_resolution` parameter.

- `list-folder-items` - List items in a Canva folder. An item can be a design, folder, or image. You can filter by item type and sort the results. Use the continuation token to get the next page of results if needed.
- `list-replies` - Get a list of replies for a specific comment on a Canva design. Comments can contain multiple replies from different users. These replies help teams collaborate by allowing discussion on a specific comment. You need to provide the design ID and comment ID. You can find the design ID using the `search-designs` tool and the comment ID using the `list-comments` tool. Use the continuation token to get the next page of results, if needed.
- `move-item-to-folder` - Move items (designs, folders, images) to a specified Canva folder
- `reply-to-comment` - Reply to an existing comment on a Canva design. You need to provide the design ID, comment ID, and your reply message. The reply will be added to the specified comment and visible to all users with access to the design.
- `resize-design` - Resize a Canva design to a preset or custom size. The tool will provide a summary of the new resized design, including its metadata.
- `search-designs` - Search docs, presentations, videos, whiteboards, sheets, and other designs in Canva. Use 'query' parameter to search by title or content. If 'query' is used, 'sortBy' must be set to 'relevance'. Filter by 'any' ownership unless specified. Sort by relevance unless specified. Use the continuation token to get the next page of results, if needed.
- `search-folders` - Search the user's folders and folders shared with the user based on folder names and tags. Returns a list of matching folders with pagination support.
- `upload-asset-from-url` - Upload an asset (e.g. an image, a video) from a URL into Canva. If the API call returns "Missing scopes: [asset:write]", you should ask the user to disconnect and reconnect their connector. This will generate a new access token with the required scope for this tool.

Setting up Canva integration

Canva integration uses MCP server connectivity to provide action capabilities. For detailed setup instructions, see [the section called "MCP integration"](#).

You'll need:

- Canva account with appropriate permissions

Compatibility

Canva integration supports:

- **Chat Agents:** Yes
- **Flows:** Yes
- **Knowledge Base:** No

Cisco Webex Meetings integration

With the Cisco Webex Meetings connector, you can schedule and manage meetings, retrieve recordings, and access participant information directly in Amazon Quick through natural language.

Amazon Quick supports multiple authentication methods for Cisco Webex Meetings. Choose the method that best fits your organization's security requirements.

- **Custom OAuth app** – Uses a customer-managed application registered in the Webex Developer portal. This option gives your organization full control over the OAuth configuration.
- **Service-to-Service OAuth** – Uses client credentials for server-to-server authentication without user interaction. Suitable for automated workflows.
- **API Key** – Uses an API key for authentication.

For more information about the authentication methods that Amazon Quick supports, see [the section called “Authentication methods”](#).

Before you begin

Make sure that you have the following before you set up the integration.

- An active Cisco Webex account with access to Webex Meetings.
- For **Custom OAuth app** or **Service-to-Service OAuth**: Access to the [Webex Developer portal](#) on the Cisco website to create an integration and obtain client credentials.
- For **API Key**: A Webex personal access token or API key with the required scopes.
- For Amazon Quick subscription requirements, see [the section called “Set up integrations in the console”](#).

Configuring Cisco Webex

For Custom OAuth app or Service-to-Service OAuth authentication, create an integration in the [Webex Developer portal](#) on the Cisco website and add the Amazon Quick callback URL `https://{region}.quicksight.aws.amazon.com/sn/oauthcallback` as a redirect URI. Replace `{region}` with your AWS Region (for example, `us-east-1`). For step-by-step instructions, see [Webex Meetings MCP Server](#) in the Webex developer documentation. Record the Client ID and Client Secret — you need them when you configure Amazon Quick. For API Key authentication, generate a Webex personal access token or API key with the required scopes.

Setting up the connector in Amazon Quick

1. In the Amazon Quick console, choose **Connectors**.
2. Choose the **Create for your team** tab.
3. Find and choose **Cisco Webex Meetings**.
4. Enter a **Name** for your connector. Optionally, choose **+ Add Description** to add a description.
5. For **Connection type**, choose **Public network**.
6. For **OAuth Configuration**, choose one of the following authentication methods and configure the required fields.
 - a. For **Custom OAuth app**, configure the following fields:
 - **Base URL** – The Webex Meetings MCP base URL. Default: `https://mcp.webexapis.com/mcp/webex-meeting`
 - **Client ID** – The client ID from your Webex Developer portal integration.
 - **Public OAuth client** (Optional) – Select this option if your Webex integration is configured as a public client (no client secret).
 - **Client secret** – The client secret from your Webex Developer portal integration.
 - **Token URL** – The token endpoint. Default: `https://webexapis.com/v1/access_token`
 - **Authorization URL** – The authorization endpoint. Default: `https://webexapis.com/v1/authorize`
 - **Redirect URL** – Pre-filled with the Amazon Quick callback URL.
 - b. For **Service-to-Service OAuth**, configure the following fields:

- **Base URL** – The Webex Meetings MCP base URL. Default: `https://mcp.webexapis.com/mcp/webex-meeting`
 - **Client ID** – The client ID from your Webex Service App.
 - **Client secret** – The client secret from your Webex Service App.
 - **Token URL** – The token endpoint. Default: `https://webexapis.com/v1/access_token`
- c. For **API Key**, configure the following fields:
- **Base URL** – The Webex Meetings MCP base URL. Default: `https://mcp.webexapis.com/mcp/webex-meeting`
 - **API Key** – The Webex API key or personal access token.
 - **Email (Optional)** – The email address associated with the API key.
7. Choose **Next**.
8. If you chose **Custom OAuth app**, a Webex authorization window opens. Review the requested permissions and choose **Accept**.
9. On the **Review** page, review the available actions for the connector. Choose **Next**.
10. On the **Publish** page, choose who can access the connector. You can enable access for everyone in your organization or search for specific teams or groups.
11. Choose **Publish**.

Available actions

After you set up the connector, the actions exposed by Cisco Webex Meetings are available. To see the current set of actions for your connector, go to the connector's **Available actions** view in the Amazon Quick console.

Managing and troubleshooting

To edit, share, or delete your connector, see [the section called “Managing existing integrations”](#).

Authentication issues

- **Sign-in fails (Custom OAuth app)** – Verify that your Cisco Webex account is active and that you can sign in to Webex directly. Confirm that the redirect URI in your Webex Developer portal integration matches the Amazon Quick callback URL.

- **Invalid client credentials** – Verify that the Client ID and Client secret match the values in your Webex Developer portal integration or Service App.
- **API Key authentication fails** – Verify that the API key has not been revoked and that it has the required scopes for Webex Meetings operations.

Coding agents

In addition to cloud connectors, you can add a coding agent, an external AI coding tool such as Kiro or Claude Code, that Amazon Quick can hand development work to. Amazon Quick communicates with the coding agent by using the Agent Client Protocol (ACP). Like other local connectors, a coding agent runs on your own machine, so it is available only when that machine is on and connected. For an overview of cloud connectors, see [Connectors](#). For information about local connectors, see [Local connectors](#).

To add a coding agent

1. Open the Desktop App, go to **Customize**, and choose the **Connectors** tab.
2. Choose **Create**, and then choose **Coding agent**.
3. Configure the connection settings, and choose **Save**.

Each coding agent shows its name, a description, and a toggle to enable or disable it. After you configure a coding agent, you can delegate tasks to it from chat. For example, you can ask Amazon Quick to use Kiro to refactor a module, or ask Claude Code to write tests for a function. Amazon Quick dispatches the work to the coding agent and reports the results to you.

Cisco Webex Video Messaging integration

With the Cisco Webex Video Messaging connector, you can manage video updates, retrieve transcripts, and access video metadata from Cisco Webex (formerly known as Vidcast) directly in Amazon Quick through natural language.

Amazon Quick supports multiple authentication methods for Cisco Webex Video Messaging. Choose the method that best fits your organization's security requirements.

- **Custom OAuth app** – Uses a customer-managed application registered in the Webex Developer portal. This option gives your organization full control over the OAuth configuration.

- **Service-to-Service OAuth** – Uses client credentials for server-to-server authentication without user interaction. Suitable for automated workflows.
- **API Key** – Uses an API key for authentication.

For more information about the authentication methods that Amazon Quick supports, see [the section called “Authentication methods”](#).

Before you begin

Make sure that you have the following before you set up the integration.

- An active Cisco Webex account with access to Webex Video Messaging (Vidcast).
- For **Custom OAuth app** or **Service-to-Service OAuth**: Access to the [Webex Developer portal](#) on the Cisco website to create an integration and obtain client credentials.
- For **API Key**: A Webex personal access token or API key with the required scopes.
- For Amazon Quick subscription requirements, see [the section called “Set up integrations in the console”](#).

Configuring Cisco Webex

For Custom OAuth app or Service-to-Service OAuth authentication, create an integration in the [Webex Developer portal](#) on the Cisco website and add the Amazon Quick callback URL `https://{region}.quicksight.aws.amazon.com/sn/oauthcallback` as a redirect URI. Replace `{region}` with your AWS Region (for example, `us-east-1`). For step-by-step instructions, see [Vidcast MCP Server](#) in the Webex developer documentation. Record the Client ID and Client Secret — you need them when you configure Amazon Quick. For API Key authentication, generate a Webex personal access token or API key with the required scopes.

Setting up the connector in Amazon Quick

1. In the Amazon Quick console, choose **Connectors**.
2. Choose the **Create for your team** tab.
3. Find and choose **Cisco Webex Video Messaging**.
4. Enter a **Name** for your connector. Optionally, choose **+ Add Description** to add a description.
5. For **Connection type**, choose **Public network**.

6. For **OAuth Configuration**, choose one of the following authentication methods and configure the required fields.
 - a. For **Custom OAuth app**, configure the following fields:
 - **Base URL** – The Webex Video Messaging MCP base URL. Default: `https://mcp.webexapis.com/mcp/vidcast`
 - **Client ID** – The client ID from your Webex Developer portal integration.
 - **Public OAuth client** (Optional) – Select this option if your Webex integration is configured as a public client (no client secret).
 - **Client secret** – The client secret from your Webex Developer portal integration.
 - **Token URL** – The token endpoint. Default: `https://webexapis.com/v1/access_token`
 - **Authorization URL** – The authorization endpoint. Default: `https://webexapis.com/v1/authorize`
 - **Redirect URL** – Pre-filled with the Amazon Quick callback URL.
 - b. For **Service-to-Service OAuth**, configure the following fields:
 - **Base URL** – The Webex Video Messaging MCP base URL. Default: `https://mcp.webexapis.com/mcp/vidcast`
 - **Client ID** – The client ID from your Webex Service App.
 - **Client secret** – The client secret from your Webex Service App.
 - **Token URL** – The token endpoint. Default: `https://webexapis.com/v1/access_token`
 - c. For **API Key**, configure the following fields:
 - **Base URL** – The Webex Video Messaging MCP base URL. Default: `https://mcp.webexapis.com/mcp/vidcast`
 - **API Key** – The Webex API key or personal access token.
 - **Email** (Optional) – The email address associated with the API key.
7. Choose **Next**.
8. If you chose **Custom OAuth app**, a Webex authorization window opens. Review the requested permissions and choose **Accept**.
9. On the **Review** page, review the available actions for the connector. Choose **Next**.

10. On the **Publish** page, choose who can access the connector. You can enable access for everyone in your organization or search for specific teams or groups.
11. Choose **Publish**.

Available actions

After you set up the connector, the actions exposed by Cisco Webex Video Messaging are available. To see the current set of actions for your connector, go to the connector's **Available actions** view in the Amazon Quick console.

Managing and troubleshooting

To edit, share, or delete your connector, see [the section called “Managing existing integrations”](#).

Authentication issues

- **Sign-in fails (Custom OAuth app)** – Verify that your Cisco Webex account is active and that you can sign in to Webex directly. Confirm that the redirect URI in your Webex Developer portal integration matches the Amazon Quick callback URL.
- **Invalid client credentials** – Verify that the Client ID and Client secret match the values in your Webex Developer portal integration or Service App.
- **API Key authentication fails** – Verify that the API key has not been revoked and that it has the required scopes for Webex Video Messaging operations.

Dropbox integration

With the Dropbox connector, you can access the Dropbox platform directly in Amazon Quick through natural language. You can upload files, manage folder structures, generate sharing links, and search for content without leaving Amazon Quick.

Amazon Quick supports multiple authentication methods for Dropbox. Choose the method that best fits your organization's security requirements.

- **Default OAuth app** – Uses an AWS-managed OAuth application. No additional credentials are needed. You authenticate directly with your Dropbox account.
- **Custom OAuth app** – Uses a customer-managed OAuth application that is registered in Dropbox. This option gives your organization full control over the OAuth configuration.

- **API Key** – Uses a Dropbox access token for authentication. This method is suitable for individual use or testing.

For more information about the authentication methods that Amazon Quick supports, see [the section called “Authentication methods”](#).

Before you begin

Make sure that you have the following before you set up the integration.

- An active Dropbox account with access to the files and folders that you want to use.
- For **Custom OAuth app**: Access to the [Dropbox App Console](#) on the Dropbox website to create an app.
- For **API Key**: An access token that is generated from the [Dropbox App Console](#) on the Dropbox website.
- For Amazon Quick subscription requirements, see [the section called “Set up integrations in the console”](#).

Configuring Dropbox

If you are using **Default OAuth app** authentication, skip this section and go to [the section called “Setting up the connector in Amazon Quick”](#).

For Custom OAuth app or API Key authentication, complete the applicable steps in Dropbox before you configure Amazon Quick.

Create a Dropbox app for Custom OAuth app

Create an app in the Dropbox App Console to get the client credentials that you need for Amazon Quick.

1. Sign in to your Dropbox account and go to the [Dropbox App Console](#) on the Dropbox website.
2. Choose **Create app**.
3. For **Choose an API**, select **Scoped access**.
4. Choose the type of access that you need:
 - **App folder** – Provides access only to a specific folder.
 - **Full Dropbox** – Provides access to all files and folders.

5. Enter an app name and choose **Create app**.
6. On the app settings page, record the following values. You need them when you configure Amazon Quick.
 - **App key** – This is your Client ID.
 - **App secret** – This is your Client secret.
7. Under **OAuth 2**, in the **Redirect URIs** section, add the Amazon Quick callback URL:
`https://{region}.quicksight.aws.amazon.com/sn/oauthcallback`
8. Choose the **Permissions** tab and select the scopes that your integration requires. For the recommended scopes, see [the section called “Recommended scopes”](#).

Generate an access token for API Key

Generate an access token from the Dropbox App Console for API Key authentication.

1. Sign in to your Dropbox account and go to the [Dropbox App Console](#) on the Dropbox website.
2. Choose your app, or create a new one if needed.
3. Under **OAuth 2**, choose **Generate** to create an access token.
4. Copy and securely store the generated token. You need it when you configure Amazon Quick.

Recommended scopes

When you create a Dropbox app, configure the following scopes on the **Permissions** tab based on the actions that you want to use.

Dropbox recommended scopes

Scope	Description
account_info.read	Reads account information.
account_info.write	Updates account information.
contacts.read	Reads contact information.
contacts.write	Updates contact information.
events.read	Reads event log entries.

Scope	Description
<code>events.write</code>	Creates event log entries.
<code>file_requests.read</code>	Reads file requests.
<code>file_requests.write</code>	Creates and manages file requests.
<code>files.content.read</code>	Reads file content.
<code>files.content.write</code>	Creates, modifies, and deletes files.
<code>files.metadata.read</code>	Reads file and folder metadata.
<code>files.metadata.write</code>	Edits file and folder metadata.
<code>files.permanent_delete</code>	Permanently deletes files.
<code>files.team_metadata.read</code>	Reads team file metadata.
<code>files.team_metadata.write</code>	Edits team file metadata.
<code>groups.read</code>	Reads group information.
<code>groups.write</code>	Manages groups.
<code>members.delete</code>	Removes team members.
<code>members.read</code>	Reads team member information.
<code>members.write</code>	Manages team members.
<code>sessions.list</code>	Lists active sessions.
<code>sessions.modify</code>	Modifies active sessions.
<code>sharing.read</code>	Reads sharing settings and shared links.
<code>sharing.write</code>	Creates and manages shared links and folder sharing.
<code>team_data.content.read</code>	Reads team content.

Scope	Description
team_data.content.write	Manages team content.
team_data.governance.read	Reads team data governance settings.
team_data.governance.write	Manages team data governance settings.
team_data.team_space	Accesses team space data.
team_info.read	Reads team information.
team_info.write	Updates team information.

Setting up the connector in Amazon Quick

Connect from the Available tab

If you want to use Default OAuth app authentication, you can connect directly from the **Available** tab without additional configuration.

1. In the Amazon Quick console, choose **Connectors**.
2. On the **Available** tab, find **Dropbox** and choose **Connect**.
3. Complete the Dropbox sign-in flow and grant the requested permissions.

To configure a connector with Custom OAuth app or API Key instead, use the **Create for your team** tab as described below.

Create from the Create for your team tab

After you complete any required Dropbox configuration, create the connector in Amazon Quick.

1. In the Amazon Quick console, choose **Connectors**.
2. Choose the **Create for your team** tab.
3. Find and choose **Dropbox**.

 Note

If a Dropbox connector already exists, a dialog appears with your existing connectors. To use an existing connector, choose it. To create a new one, choose **No, create new**.

4. Enter a **Name** for your connector. Optionally, choose **+ Add Description** to add a description.
5. For **Connection type**, choose **Public network**.
6. For **OAuth Configuration**, choose one of the following authentication methods and configure the required fields.

- a. For **Default OAuth app**:

No additional credentials are needed. Choose **Next** to continue.

- b. For **Custom OAuth app**, configure the following fields:

- **Base URL** (Optional) – The Dropbox API base URL. Example: `https://api.dropboxapi.com`
- **Client ID** – The app key from your Dropbox app.
- **Client secret** – The app secret from your Dropbox app.
- **Token URL** – The token endpoint. Example: `https://api.dropboxapi.com/oauth2/token`
- **Authorization URL** – The authorization endpoint. Example: `https://www.dropbox.com/oauth2/authorize`
- **Redirect URL** – Pre-filled with the Amazon Quick callback URL.

- c. For **API Key**, configure the following fields:

- **Base URL** (Optional) – The Dropbox API base URL. Example: `https://api.dropboxapi.com`
- **API Key** – Your Dropbox access token.
- **Email** (Optional) – The email address that is associated with your Dropbox account.

7. Choose **Next**.
8. If you chose **Default OAuth app** or **Custom OAuth app**, a Dropbox authorization window opens. Review the requested permissions and choose **Allow**.
9. On the **Review** page, review the available actions for the connector. Choose **Next**.

10. On the **Publish** page, choose who can access the connector. You can turn on access for everyone in your organization or search for specific teams or groups.
11. Choose **Publish**.

Available actions

After you set up the connector, the following actions are available.

Dropbox available actions

Category	Action	Description
Files	List Files	Lists files in a folder.
Files	Get File	Reads the content of a file.
Files	Import File	Uploads a file to Dropbox.
Files	Search Files	Searches for files or folders by name or content.
Files	Update File Location	Moves a file or folder to a different location.
Files	Get Temporary Link	Generates a temporary download link for a file.
Files	Get Metadata	Retrieves metadata for a file or folder.
Files	Save File From URL	Saves a file to Dropbox from a URL.
Folders	Create Folder	Creates a new folder.
Folders	List Folders	List of folders user has access to.
Batch operations	Get Batch Move Status	Checks the status of a batch move operation.
Sharing	Create Shared Link	Creates a shared link for a file or folder.
Sharing	List Shared Links	Lists shared links for a file or folder.
Sharing	Update Folder Sharing	Shares a folder with other users.

Category	Action	Description
Sharing	List Received Files	Lists files that other users have shared with you.
File requests	Create File Request	Creates a file request for others to upload files to.
Account	Get User	Retrieves information about the authenticated user.
Account	Get Space Usage	Retrieves storage space usage for the account.

 Note

The actions that you can use depend on the files and folders that are accessible to the authenticated user.

Managing and troubleshooting

To edit, share, or delete your connector, see [the section called “Managing existing integrations”](#).

Authentication issues

- **Sign-in fails (Default OAuth app or Custom OAuth app)** – Verify that your Dropbox account is active and that you can sign in to [the Dropbox website](#) directly. For Custom OAuth app, confirm that the redirect URL in your Dropbox app matches the Amazon Quick callback URL.
- **Invalid client credentials (Custom OAuth app)** – Verify that the Client ID and Client secret match the app key and app secret values in your Dropbox app. If you regenerated the app secret, update the value in Amazon Quick.
- **API Key rejected** – Verify that your access token is active. Tokens can expire or be revoked in the Dropbox App Console.
- **Insufficient permissions** – Verify that your Dropbox app has the required permission scopes for the actions that you want to use.

Dun & Bradstreet integration

With the Dun & Bradstreet connector, you can look up company profiles, credit risk data, and business intelligence directly in Amazon Quick through natural language.

Amazon Quick supports multiple authentication methods for Dun & Bradstreet. Choose the method that best fits your organization's security requirements.

- **Custom OAuth app** – Uses a customer-managed OAuth application registered in the Dun & Bradstreet developer portal. This option gives your organization full control over the OAuth configuration.
- **Service-to-Service OAuth** – Uses client credentials for server-to-server authentication without user interaction. Suitable for automated workflows.

For more information about the authentication methods that Amazon Quick supports, see [the section called “Authentication methods”](#).

Before you begin

Make sure that you have the following before you set up the integration.

- An active Dun & Bradstreet D&B Direct+ account with access to the data products you want to query.
- For **Custom OAuth app** or **Service-to-Service OAuth**: OAuth credentials from your Dun & Bradstreet developer portal.
- For Amazon Quick subscription requirements, see [the section called “Set up integrations in the console”](#).

Configuring Dun & Bradstreet

For Custom OAuth app or Service-to-Service OAuth authentication, register an OAuth app in the Dun & Bradstreet developer portal and add the Amazon Quick callback URL `https://{region}.quicksight.aws.amazon.com/sn/oauthcallback` as a redirect URI. Replace `{region}` with your AWS Region (for example, `us-east-1`). For authentication details, see [Authentication](#) in the [D&B Direct+ documentation](#). Record the Client ID and Client Secret — you need them when you configure Amazon Quick.

Setting up the connector in Amazon Quick

1. In the Amazon Quick console, choose **Connectors**.
2. Choose the **Create for your team** tab.
3. Find and choose **Dun & Bradstreet**.
4. Enter a **Name** for your connector. Optionally, choose **+ Add Description** to add a description.
5. For **Connection type**, choose **Public network**.
6. For **OAuth Configuration**, choose one of the following authentication methods and configure the required fields.
 - a. For **Custom OAuth app**, configure the following fields:
 - **Base URL** – The Dun & Bradstreet MCP base URL. Default: `https://plus.dnb.com/v2/mcp`
 - **Client ID** – The client ID from your Dun & Bradstreet OAuth app.
 - **Public OAuth client** (Optional) – Select this option if your Dun & Bradstreet OAuth app is configured as a public client (no client secret).
 - **Client secret** – The client secret from your Dun & Bradstreet OAuth app.
 - **Token URL** – The token endpoint. Default: `https://plus.dnb.com/v2/mcp/token`
 - **Authorization URL** – The authorization endpoint. Default: `https://plus.dnb.com/v2/mcp/authorize`
 - **Redirect URL** – Pre-filled with the Amazon Quick callback URL.
 - b. For **Service-to-Service OAuth**, configure the following fields:
 - **Base URL** – The Dun & Bradstreet MCP base URL. Default: `https://plus.dnb.com/v2/mcp`
 - **Client ID** – The client ID from your Dun & Bradstreet OAuth app.
 - **Client secret** – The client secret from your Dun & Bradstreet OAuth app.
 - **Token URL** – The token endpoint. Default: `https://plus.dnb.com/v2/mcp/token`
7. Choose **Next**.
8. If you chose **Custom OAuth app**, a Dun & Bradstreet authorization window opens. Review the requested permissions and choose **Allow**.
9. On the **Review** page, review the available actions for the connector. Choose **Next**.

10. On the **Publish** page, choose who can access the connector. You can enable access for everyone in your organization or search for specific teams or groups.
11. Choose **Publish**.

Available actions

After you set up the connector, the actions exposed by Dun & Bradstreet are available. To see the current set of actions for your connector, go to the connector's **Available actions** view in the Amazon Quick console.

Managing and troubleshooting

To edit, share, or delete your connector, see [the section called “Managing existing integrations”](#).

Authentication issues

- **Sign-in fails (Custom OAuth app)** – Verify that your Dun & Bradstreet account is active. Confirm that the redirect URI in your Dun & Bradstreet OAuth app matches the Amazon Quick callback URL.
- **Invalid client credentials** – Verify that the Client ID and Client secret match the values in your Dun & Bradstreet OAuth app.

GitHub integration

With GitHub integration in Amazon Quick, you can manage repositories, create and review pull requests, track issues, and collaborate on code through MCP server connectivity. This integration provides action capabilities for development workflow operations.

What you can do

GitHub integration provides connector capabilities through MCP server connectivity:

- Create and manage repositories
- Create, review, and merge pull requests
- Track and manage issues
- Manage branches and commits
- Review code and provide feedback
- Manage project boards and milestones

Available tools

The GitHub MCP server typically provides these tools:

- `create_repository` - Create new repositories
- `get_repository` - Get repository information
- `list_issues` - List repository issues
- `create_issue` - Create new issues
- `create_pull_request` - Create pull requests
- `list_pull_requests` - List pull requests
- `get_file_contents` - Read file contents
- `create_commit` - Create commits
- `search_repositories` - Search for repositories

Note

The specific tools and capabilities available through this MCP server may change over time. For the most current information about supported tools, features, and implementation details, check the official GitHub documentation and MCP server repository.

Setting up GitHub integration

GitHub integration uses MCP server connectivity to provide action capabilities. For detailed setup instructions, see [the section called “MCP integration”](#).

You'll need:

- GitHub account with appropriate repository permissions
- Personal access token or GitHub App credentials

Compatibility

GitHub integration supports:

- **Chat Agents:** Yes

- **Flows:** Yes
- **Knowledge Base:** No

Gmail integration

With the Gmail connector, you can access Gmail directly in Amazon Quick through natural language. You can read and send emails, manage labels and drafts, organize threads, and search contacts without leaving Amazon Quick.

Amazon Quick supports multiple authentication methods for Gmail. Choose the method that best fits your organization's security requirements.

- **Default OAuth app** – Uses an AWS-managed OAuth application. No additional credentials are needed. You authenticate directly with your Google account.
- **Custom OAuth app** – Uses a customer-managed OAuth client created in the Google Cloud Console. This option gives your organization full control over the OAuth configuration.

For more information about the authentication methods that Amazon Quick supports, see [the section called “Authentication methods”](#).

Before you begin

Make sure that you have the following before you set up the integration.

- A Google account with access to Gmail.
- For **Custom OAuth app**: Access to the [Google Cloud Console](#) on the Google website with permissions to create OAuth clients.
- For Amazon Quick subscription requirements, see [the section called “Set up integrations in the console”](#).

Configuring Google Cloud

If you use **Default OAuth app** authentication, skip this section and see [the section called “Setting up the connector in Amazon Quick”](#).

For Custom OAuth app authentication, complete the following steps in the Google Cloud Console before you configure Amazon Quick. When you enable the API in step 3, search for and enable the **Gmail API**.

Create an OAuth client in Google Cloud Console

Create an OAuth client in the Google Cloud Console to get the client credentials that you need for Amazon Quick. For more information, see [Using OAuth 2.0 to Access Google APIs](#) on the Google website.

1. Sign in to the [Google Cloud Console](#) on the Google website.
2. Create a new project or select an existing project.
3. In the left navigation pane, choose **APIs & Services**, then choose **Library**. Search for the API that your integration requires and choose **Enable**.
4. Choose **OAuth consent screen** and choose **Get Started**.
5. Configure the consent screen:
 - Enter an **App Name** and select a **User support email**.
 - For **Audience**, choose **Internal** (your organization only) or **External** (any Google user).
 - Add developer contact details and choose **Create**.
6. Choose **Create OAuth client**.
7. Configure the client:
 - For **Application type**, choose **Web application**.
 - Enter a **Name** for your client.
 - Under **Authorized redirect URIs**, add the Amazon Quick callback URL:
`https://{region}.quicksight.aws.amazon.com/sn/oauthcallback`
8. Choose **Create**.
9. Record the following values. You need them when you configure Amazon Quick.
 - **Client ID**
 - **Client secret**

Recommended scopes

The following scopes are requested when you connect to Gmail.

Gmail recommended scopes

Scope	Description
https://www.googleapis.com/auth/gmail.modify	Reads, sends, and manages emails and labels.
https://www.googleapis.com/auth/gmail.send	Sends emails on behalf of the user.
https://www.googleapis.com/auth/contacts.readonly	Reads contact information.
openid	Authenticates the user's identity.
email	Reads the user's email address.

Note

The `openid` and `email` scopes are automatically included in the OAuth consent flow for user authentication. They are not required for Service-to-Service OAuth.

Setting up the connector in Amazon Quick

Connect from the Available tab

If you want to use Default OAuth app authentication, you can connect directly from the **Available** tab without additional configuration.

1. In the Amazon Quick console, choose **Connectors**.
2. On the **Available** tab, find **Gmail** and choose **Connect**.
3. Complete the Google sign-in flow and grant the requested permissions.

To configure a connector with Custom OAuth app instead, use the **Create for your team** tab as described below.

Create from the Create for your team tab

After you complete any required Google Cloud configuration, create the connector in Amazon Quick.

1. In the Amazon Quick console, choose **Connectors**.
2. Choose the **Create for your team** tab.
3. Find and choose **Gmail**.

Note

If a Gmail connector already exists, a dialog appears with your existing connectors. To use an existing connector, choose it. To create a new one, choose **No, create new**.

4. Enter a **Name** for your connector. Optionally, choose **+ Add Description** to add a description.
5. For **Connection type**, choose **Public network**.
6. For **OAuth Configuration**, choose one of the following authentication methods and configure the required fields.
 - a. For **Default OAuth app**:

No additional credentials are needed. Choose **Next** to continue.
 - b. For **Custom OAuth app**, configure the following fields:
 - **Base URL** (Optional) – The Gmail API base URL. For example: `https://gmail.googleapis.com`
 - **Client ID** – The client ID from your Google Cloud OAuth client.
 - **Client secret** – The client secret from your Google Cloud OAuth client.
 - **Token URL** – The token endpoint. For example: `https://oauth2.googleapis.com/token`
 - **Authorization URL** – The authorization endpoint. For example: `https://accounts.google.com/o/oauth2/v2/auth`
 - **Redirect URL** – Pre-filled with the Amazon Quick callback URL.
7. Choose **Next**.
8. If you chose **Default OAuth app** or **Custom OAuth app**, a Google authorization window opens. Review the requested permissions and choose **Allow**.

9. On the **Review** page, review the available actions for the connector. Choose **Next**.
10. On the **Publish** page, choose who can access the connector. You can enable access for everyone in your organization or search for specific teams or groups.
11. Choose **Publish**.

Available actions

After you set up the connector, the following actions are available.

Gmail available actions

Category	Action	Description
Messages	Get Emails	Retrieves emails from the mailbox.
Messages	Get Message By Thread Id	Retrieves a message by its thread ID.
Messages	Get Message By Message Id	Retrieves a message by its message ID.
Messages	Send Email	Sends a new email message.
Messages	Send Thread Reply	Replies to an existing email thread.
Messages	Forward Message	Forwards an email to other recipients.
Messages	Get Attachment	Retrieves an email attachment.
Messages	Batch Update Message	Modifies labels on multiple messages in a single operation.
Threads	List Threads	Lists email threads in the mailbox.
Threads	Update Thread Labels	Modifies labels on a thread.
Trash	Delete Email	Moves an email to the trash.
Trash	Delete Thread	Moves a thread to the trash.
Trash	Restore Message	Restores an email from the trash.
Trash	Restore Thread	Restores a thread from the trash.

Category	Action	Description
Labels	List Labels	Lists all labels in the mailbox.
Labels	Get Label	Retrieves details for a specific label.
Labels	Create Label	Creates a new label.
Labels	Update Label	Updates a label's properties.
Labels	Tag Email	Adds a label to an email.
Drafts	List Drafts	Lists email drafts.
Drafts	Get Draft	Retrieves a specific draft.
Drafts	Create Draft	Creates a new email draft.
Drafts	Update Draft	Updates an existing draft.
Drafts	Send Draft	Sends an existing draft.
Contacts	Search Contacts	Searches for contacts by name or email.
Contacts	Get Contacts	Retrieves contact information.
History	List History	Lists mailbox change history.
Settings	Get Profile	Retrieves the authenticated user's Gmail profile.
Settings	Get Send As Alias	Retrieves a specific send-as alias configuration.
Settings	List Send As Aliases	Lists all send-as aliases for the account.

 Note

The actions that you can use depend on the data accessible to the authenticated user.

Manage and troubleshoot

To edit, share, or delete your connector, see [the section called “Managing existing integrations”](#).

Common Google authentication issues

- **Sign-in fails (Default OAuth app or Custom OAuth app)** – Verify that your Google account is active and that you can sign in to [the Google website](#) directly. For Custom OAuth app, confirm that the redirect URI in your Google Cloud OAuth client matches the Amazon Quick callback URL.
- **App blocked by administrator** – If your Google Workspace administrator restricts third-party app access, you might see an error when you attempt to sign in. Contact your Google Workspace administrator to allow the Amazon Quick app.
- **Invalid client credentials (Custom OAuth app)** – Verify that the Client ID and Client secret match the values in your Google Cloud OAuth client.
- **Authentication popup fails** – Verify that your browser allows popups from the Amazon Quick console domain. Try using a different browser or clearing your browser cache.
- **Permissions revoked** – If you previously revoked Amazon Quick access from your Google Account permissions settings, you need to re-authenticate by editing the connector and signing in again.
- **Google API rate limiting** – Google might limit requests during high usage periods. If actions fail, retry after a few minutes.

Gmail-specific issues

- **Gmail API not enabled** – Verify that the Gmail API is enabled in your Google Cloud project under **APIs & Services, Library**.

Google Analytics integration

With the Google Analytics connector, you can access Google Analytics directly in Amazon Quick through natural language. You can run reports, retrieve audience data, monitor key events, and access real-time metrics without leaving Amazon Quick.

Amazon Quick supports multiple authentication methods for Google Analytics. Choose the method that best fits your organization's security requirements.

- **Default OAuth app** – Uses an AWS-managed OAuth application. No additional credentials are needed. You authenticate directly with your Google account.
- **Custom OAuth app** – Uses a customer-managed OAuth client created in the Google Cloud Console. This option gives your organization full control over the OAuth configuration.

For more information about the authentication methods that Amazon Quick supports, see [the section called “Authentication methods”](#).

Before you begin

Make sure that you have the following before you set up the integration.

- A Google account with access to Google Analytics properties and data that you want to use.
- For **Custom OAuth app**: Access to the [Google Cloud Console](#) on the Google website with permissions to create OAuth clients.
- For Amazon Quick subscription requirements, see [the section called “Set up integrations in the console”](#).

Configuring Google Cloud

If you are using **Default OAuth app** authentication, skip this section and go to [the section called “Setting up the connector in Amazon Quick”](#).

For Custom OAuth app authentication, complete the following steps in the Google Cloud Console before you configure Amazon Quick. When you enable the API in step 3, search for and enable the **Google Analytics Data API**.

Create an OAuth client in Google Cloud Console

Create an OAuth client in the Google Cloud Console to get the client credentials that you need for Amazon Quick. For more information, see [Using OAuth 2.0 to Access Google APIs](#) on the Google website.

1. Sign in to the [Google Cloud Console](#) on the Google website.
2. Create a new project or select an existing project.
3. In the left navigation pane, choose **APIs & Services**, then choose **Library**. Search for the API that your integration requires and choose **Enable**.

4. Choose **OAuth consent screen** and choose **Get Started**.
5. Configure the consent screen:
 - Enter an **App Name** and select a **User support email**.
 - For **Audience**, choose **Internal** (your organization only) or **External** (any Google user).
 - Add developer contact details and choose **Create**.
6. Choose **Create OAuth client**.
7. Configure the client:
 - For **Application type**, choose **Web application**.
 - Enter a **Name** for your client.
 - Under **Authorized redirect URIs**, add the Amazon Quick callback URL:
`https://{region}.quicksight.aws.amazon.com/sn/oauthcallback`
8. Choose **Create**.
9. Record the following values. You need them when you configure Amazon Quick.
 - **Client ID**
 - **Client secret**

Recommended scopes

The following scopes are requested when you connect to Google Analytics.

Google Analytics recommended scopes

Scope	Description
<code>https://www.googleapis.com/auth/analytics.readonly</code>	Reads Analytics data and configuration.
<code>https://www.googleapis.com/auth/analytics.edit</code>	Edits Analytics configuration and data.
<code>openid</code>	Authenticates the user's identity.
<code>email</code>	Reads the user's email address.

 Note

The `openid` and `email` scopes are automatically included in the OAuth consent flow for user authentication. They are not required for Service-to-Service OAuth.

Setting up the connector in Amazon Quick

Connect from the Available tab

If you want to use Default OAuth app authentication, you can connect directly from the **Available** tab without additional configuration.

1. In the Amazon Quick console, choose **Connectors**.
2. On the **Available** tab, find **Google Analytics** and choose **Connect**.
3. Complete the Google sign-in flow and grant the requested permissions.

To configure a connector with Custom OAuth app instead, use the **Create for your team** tab as described below.

Create from the Create for your team tab

After you complete any required Google Cloud configuration, create the connector in Amazon Quick.

1. In the Amazon Quick console, choose **Connectors**.
2. Choose the **Create for your team** tab.
3. Find and choose **Google Analytics**.

 Note

If a Google Analytics connector already exists, a dialog appears with your existing connectors. To use an existing connector, choose it. To create a new one, choose **No, create new**.

4. Enter a **Name** for your connector. Optionally, choose **+ Add Description** to add a description.
5. For **Connection type**, choose **Public network**.

6. For **OAuth Configuration**, choose one of the following authentication methods and configure the required fields.
 - a. For **Default OAuth app**:

No additional credentials are needed. Choose **Next** to continue.
 - b. For **Custom OAuth app**, configure the following fields:
 - **Base URL** (Optional) – The Google Analytics API base URL. Example: `https://analyticsdata.googleapis.com`
 - **Client ID** – The client ID from your Google Cloud OAuth client.
 - **Client secret** – The client secret from your Google Cloud OAuth client.
 - **Token URL** – The token endpoint. Example: `https://oauth2.googleapis.com/token`
 - **Authorization URL** – The authorization endpoint. Example: `https://accounts.google.com/o/oauth2/v2/auth`
 - **Redirect URL** – Pre-filled with the Amazon Quick callback URL.
7. Choose **Next**.
8. If you chose **Default OAuth app** or **Custom OAuth app**, a Google authorization window opens. Review the requested permissions and choose **Allow**.
9. On the **Review** page, review the available actions for the connector. Choose **Next**.
10. On the **Publish** page, choose who can access the connector. You can enable access for everyone in your organization or search for specific teams or groups.
11. Choose **Publish**.

Available actions

After you set up the connector, the following actions are available.

Google Analytics available actions

Category	Action	Description
Reports	Create Report	Runs a report for a property.
Reports	Batch Create Report	Runs multiple reports in a single request.

Category	Action	Description
Reports	Create Realtime Report	Runs a real-time report for a property.
Reports	Create Funnel Report	Runs a funnel report for a property.
Audiences	List Audiences	Lists audiences for a property.
Audiences	Create Audience List	Creates an audience list for export.
Audiences	List Audience Lists	Lists audience lists for a property.
Audiences	Get Audience List	Retrieves data from an audience list.
Audiences	Create Audience Export	Creates an audience export.
Audiences	Get Audience Export	Retrieves data from an audience export.
Accounts	List Account Summaries	Lists account summaries accessible to the user.
Accounts	List Accounts	Lists all accounts accessible to the user.
Accounts	Get Account	Retrieves details for a specific account.
Properties	Get Property	Retrieves details for a specific property.
Properties	Search Properties	Searches for properties with filtering.
Properties	Get Property Quotas	Retrieves quota usage for a property.
Events	List Key Events	Lists key events (conversions) for a property.
Metadata	Get Metadata	Retrieves available dimensions and metrics for reporting.
Metadata	Get Compatibility	Checks compatibility of dimensions and metrics for a report.

Category	Action	Description
Annotations	List Annotations	Lists reporting data annotations for a property.

Note

The actions that you can use depend on the Analytics properties and accounts accessible to the authenticated user.

Managing and troubleshooting

To edit, share, or delete your connector, see [the section called “Managing existing integrations”](#).

Common Google authentication issues

- **Sign-in fails (Default OAuth app or Custom OAuth app)** – Verify that your Google account is active and that you can sign in to [the Google website](#) directly. For Custom OAuth app, confirm that the redirect URI in your Google Cloud OAuth client matches the Amazon Quick callback URL.
- **App blocked by administrator** – If your Google Workspace administrator restricts third-party app access, you might see an error when you attempt to sign in. Contact your Google Workspace administrator to allow the Amazon Quick app.
- **Invalid client credentials (Custom OAuth app)** – Verify that the Client ID and Client secret match the values in your Google Cloud OAuth client.
- **Authentication popup fails** – Verify that your browser allows popups from the Amazon Quick console domain. Try using a different browser or clearing your browser cache.
- **Permissions revoked** – If you previously revoked Amazon Quick access from your Google Account permissions settings, you need to re-authenticate by editing the connector and signing in again.
- **Google API rate limiting** – Google might limit requests during high usage periods. If actions fail, retry after a few minutes.

Google Analytics-specific issues

- **Google Analytics Data API not enabled** – Verify that the Google Analytics Data API is enabled in your Google Cloud project under **APIs & Services, Library**.
- **No properties visible** – Verify that the authenticated Google account has access to the Google Analytics properties that you want to query.

Google Calendar integration

With the Google Calendar connector, you can access Google Calendar directly in Amazon Quick through natural language. You can create and update events, check availability, manage calendars, and organize attendees without leaving Amazon Quick.

Amazon Quick supports multiple authentication methods for Google Calendar. Choose the method that best fits your organization's security requirements.

- **Default OAuth app** – Uses an AWS-managed OAuth application. No additional credentials are needed. You authenticate directly with your Google account.
- **Custom OAuth app** – Uses a customer-managed OAuth client created in the Google Cloud Console. This option gives your organization full control over the OAuth configuration.

For more information about the authentication methods that Amazon Quick supports, see [the section called “Authentication methods”](#).

Before you begin

Make sure that you have the following before you set up the integration.

- A Google account with access to Google Calendar.
- For **Custom OAuth app**: Access to the [Google Cloud Console](#) on the Google website with permissions to create OAuth clients.
- For Amazon Quick subscription requirements, see [the section called “Set up integrations in the console”](#).

Configuring Google Cloud

If you are using **Default OAuth app** authentication, skip this section and go to [the section called “Setting up the connector in Amazon Quick”](#).

For Custom OAuth app authentication, complete the following steps in the Google Cloud Console before you configure Amazon Quick. When you enable the API in step 3, search for and enable the **Google Calendar API**.

Create an OAuth client in Google Cloud Console

Create an OAuth client in the Google Cloud Console to get the client credentials that you need for Amazon Quick. For more information, see [Using OAuth 2.0 to Access Google APIs](#) on the Google website.

1. Sign in to the [Google Cloud Console](#) on the Google website.
2. Create a new project or select an existing project.
3. In the left navigation pane, choose **APIs & Services**, then choose **Library**. Search for the API that your integration requires and choose **Enable**.
4. Choose **OAuth consent screen** and choose **Get Started**.
5. Configure the consent screen:
 - Enter an **App Name** and select a **User support email**.
 - For **Audience**, choose **Internal** (your organization only) or **External** (any Google user).
 - Add developer contact details and choose **Create**.
6. Choose **Create OAuth client**.
7. Configure the client:
 - For **Application type**, choose **Web application**.
 - Enter a **Name** for your client.
 - Under **Authorized redirect URIs**, add the Amazon Quick callback URL:
`https://{region}.quicksight.aws.amazon.com/sn/oauthcallback`
8. Choose **Create**.
9. Record the following values. You need them when you configure Amazon Quick.
 - **Client ID**
 - **Client secret**

Recommended scopes

The following scopes are requested when you connect to Google Calendar.

Google Calendar recommended scopes

Scope	Description
https://www.googleapis.com/auth/calendar	Reads and writes calendar events and settings.
openid	Authenticates the user's identity.
email	Reads the user's email address.

Note

The `openid` and `email` scopes are automatically included in the OAuth consent flow for user authentication. They are not required for Service-to-Service OAuth.

Setting up the connector in Amazon Quick

Connect from the Available tab

If you want to use Default OAuth app authentication, you can connect directly from the **Available** tab without additional configuration.

1. In the Amazon Quick console, choose **Connectors**.
2. On the **Available** tab, find **Google Calendar** and choose **Connect**.
3. Complete the Google sign-in flow and grant the requested permissions.

To configure a connector with Custom OAuth app instead, use the **Create for your team** tab as described below.

Create from the Create for your team tab

After you complete any required Google Cloud configuration, create the connector in Amazon Quick.

1. In the Amazon Quick console, choose **Connectors**.
2. Choose the **Create for your team** tab.
3. Find and choose **Google Calendar**.

 Note

If a Google Calendar connector already exists, a dialog appears with your existing connectors. To use an existing connector, choose it. To create a new one, choose **No, create new**.

4. Enter a **Name** for your connector. Optionally, choose **+ Add Description** to add a description.
5. For **Connection type**, choose **Public network**.
6. For **OAuth Configuration**, choose one of the following authentication methods and configure the required fields.
 - a. For **Default OAuth app**:

No additional credentials are needed. Choose **Next** to continue.
 - b. For **Custom OAuth app**, configure the following fields:
 - **Base URL** (Optional) – The Google Calendar API base URL. Example: `https://www.googleapis.com/calendar`
 - **Client ID** – The client ID from your Google Cloud OAuth client.
 - **Client secret** – The client secret from your Google Cloud OAuth client.
 - **Token URL** – The token endpoint. Example: `https://oauth2.googleapis.com/token`
 - **Authorization URL** – The authorization endpoint. Example: `https://accounts.google.com/o/oauth2/v2/auth`
 - **Redirect URL** – Pre-filled with the Amazon Quick callback URL.
7. Choose **Next**.
8. If you chose **Default OAuth app** or **Custom OAuth app**, a Google authorization window opens. Review the requested permissions and choose **Allow**.
9. On the **Review** page, review the available actions for the connector. Choose **Next**.
10. On the **Publish** page, choose who can access the connector. You can enable access for everyone in your organization or search for specific teams or groups.

11. Choose **Publish**.

Available actions

After you set up the connector, the following actions are available.

Google Calendar available actions

Category	Action	Description
Events	List Events	Lists events on a calendar.
Events	Search Events	Searches for events by keyword or criteria.
Events	Get Event	Retrieves details for a specific event.
Events	Create Event	Creates a new calendar event.
Events	Create Quick Event	Creates an event from a text string.
Events	Update Event	Updates an existing event.
Events	Update Event Calendar	Moves an event to a different calendar.
Events	Import Events	Imports events into a calendar.
Events	List Event Instances	Lists instances of a recurring event.
Events	Delete Attendee	Removes an attendee from an event.
Availability	List Free Slots	Finds available time slots across calendars.
Calendars	List Calendars	Lists calendars for the authenticated user.
Calendars	List Calendar Events	Lists events across all calendars.
Calendars	Get Calendar	Retrieves details for a specific calendar.
Calendars	Get Calendar List	Retrieves a calendar list entry.
Calendars	Update Calendar List	Updates a calendar list entry.

Category	Action	Description
Settings	List Settings	Lists calendar settings for the user.
Utilities	Get Current Date Time	Retrieves the current date and time.

Note

The actions that you can use depend on the calendars accessible to the authenticated user.

Managing and troubleshooting

To edit, share, or delete your connector, see [the section called “Managing existing integrations”](#).

Common Google authentication issues

- **Sign-in fails (Default OAuth app or Custom OAuth app)** – Verify that your Google account is active and that you can sign in to [the Google website](#) directly. For Custom OAuth app, confirm that the redirect URI in your Google Cloud OAuth client matches the Amazon Quick callback URL.
- **App blocked by administrator** – If your Google Workspace administrator restricts third-party app access, you might see an error when you attempt to sign in. Contact your Google Workspace administrator to allow the Amazon Quick app.
- **Invalid client credentials (Custom OAuth app)** – Verify that the Client ID and Client secret match the values in your Google Cloud OAuth client.
- **Authentication popup fails** – Verify that your browser allows popups from the Amazon Quick console domain. Try using a different browser or clearing your browser cache.
- **Permissions revoked** – If you previously revoked Amazon Quick access from your Google Account permissions settings, you need to re-authenticate by editing the connector and signing in again.
- **Google API rate limiting** – Google might limit requests during high usage periods. If actions fail, retry after a few minutes.

Google Calendar-specific issues

- **Google Calendar API not enabled** – Verify that the Google Calendar API is enabled in your Google Cloud project under **APIs & Services, Library**.

Google Chat integration

With the Google Chat connector, you can send messages, manage spaces, and interact with chat threads from your Google Workspace account directly in Amazon Quick through natural language.

Google Chat uses Custom OAuth app authentication. For more information about the authentication methods that Amazon Quick supports, see [the section called “Authentication methods”](#).

Before you begin

Make sure that you have the following before you set up the integration.

- An active Google Workspace account with access to the Google Chat spaces that you want to use.
- An OAuth 2.0 client registered in the Google Cloud console for your Google Workspace project, with the Google Chat API enabled.
- A Google Chat App configured in the Google Cloud console under the Google Chat API settings. The Chat App must be configured before you can use the Chat API to send messages and interact with spaces. For more information, see [the section called “Configure the Google Chat App”](#).
- For Amazon Quick subscription requirements, see [the section called “Set up integrations in the console”](#).

Configuring Google Cloud

For Custom OAuth app authentication, create an OAuth 2.0 client in the [Google Cloud console](#), enable the Google Chat API for your project, configure the OAuth consent screen with the Chat scopes that you need (for example, `chat.messages`, `chat.spaces`), and add the Amazon Quick callback URL `https://{region}.quicksight.aws.amazon.com/sn/oauthcallback` as an authorized redirect URI. Replace `{region}` with your AWS Region (for example, `us-east-1`). For step-by-step instructions, see [Google Chat API](#) in the Google Workspace developer documentation. Record the Client ID and Client Secret — you need them when you configure Amazon Quick.

Configure the Google Chat App

Before you can use the Google Chat API, you must configure a Google Chat App in the Google Cloud console. The Chat App is the identity under which your integration interacts with Google Chat spaces and threads.

1. In the [Google Cloud console](#), navigate to **APIs & Services** and confirm that the **Google Chat API** is enabled for your project.
2. Choose **Google Chat API**, then choose the **Configuration** tab.
3. Enter the following Chat App details:
 - **App name.** A display name for the Chat App (for example, "Amazon Quick Integration").
 - **Avatar URL (Optional).** A URL to an image that represents the Chat App in Google Chat.
 - **Description.** A brief description of the Chat App's purpose.
4. Under **Interactive features**, configure the features that your integration requires. If your integration only sends messages and does not need to receive events from Google Chat, you can leave interactive features turned off.
5. Under **Visibility**, choose who can install the Chat App in your Google Workspace domain. You can make the app available to specific people or to everyone in your organization.
6. Choose **Save**.

For more information about configuring Google Chat Apps, see [Google Chat app overview](#) in the Google Workspace developer documentation.

Setting up the connector in Amazon Quick

1. In the Amazon Quick console, choose **Connectors**.
2. Choose the **Create for your team** tab.
3. Find and choose **Google Chat**.
4. Enter a **Name** for your connector. Optionally, choose **+ Add Description** to add a description.
5. For **Connection type**, choose **Public network**.
6. For **OAuth Configuration**, choose **Custom OAuth app** and configure the following fields:
 - **Base URL** – The Google Chat MCP base URL.
 - **Client ID** – The OAuth client ID from your Google Cloud OAuth 2.0 client.

- **Public OAuth client** (Optional) – Select this option if your Google Cloud OAuth 2.0 client is configured as a public client (no client secret).
 - **Client secret** – The client secret from your Google Cloud OAuth 2.0 client.
 - **Token URL** – The Google OAuth token endpoint.
 - **Authorization URL** – The Google OAuth authorization endpoint.
 - **Redirect URL** – Pre-filled with the Amazon Quick callback URL.
7. Choose **Next**.
 8. A Google sign-in window opens. Review the requested permissions and choose **Allow**.
 9. On the **Review** page, review the available actions for the connector. Choose **Next**.
 10. On the **Publish** page, choose who can access the connector. You can enable access for everyone in your organization or search for specific teams or groups.
 11. Choose **Publish**.

Available actions

After you set up the connector, the actions exposed by Google Chat are available. To see the current set of actions for your connector, go to the connector's **Available actions** view in the Amazon Quick console.

Managing and troubleshooting

To edit, share, or delete your connector, see [the section called “Managing existing integrations”](#).

Authentication issues

- **Sign-in fails** – Verify that your Google Workspace account is active and that you can sign in to Google Chat directly. Confirm that the authorized redirect URI in your Google Cloud OAuth 2.0 client matches the Amazon Quick callback URL.
- **Invalid client credentials** – Verify that the Client ID and Client Secret match the values in your Google Cloud OAuth 2.0 client.
- **Insufficient permissions** – Verify that the Chat API is enabled and that your OAuth consent screen includes the required scopes for the operations you want to use.

Google Docs integration

With the Google Docs connector, you can access Google Docs directly in Amazon Quick through natural language. You can create and edit documents, read content, insert text and images, and manage document structure without leaving Amazon Quick.

Amazon Quick supports multiple authentication methods for Google Docs. Choose the method that best fits your organization's security requirements.

- **Default OAuth app** – Uses an AWS-managed OAuth application. No additional credentials are needed. You authenticate directly with your Google account.
- **Custom OAuth app** – Uses a customer-managed OAuth client created in the Google Cloud Console. This option gives your organization full control over the OAuth configuration.

For more information about the authentication methods that Amazon Quick supports, see [the section called “Authentication methods”](#).

Before you begin

Make sure that you have the following before you set up the integration.

- A Google account with access to Google Docs.
- For **Custom OAuth app**: Access to the [Google Cloud Console](#) on the Google website with permissions to create OAuth clients.
- For Amazon Quick subscription requirements, see [the section called “Set up integrations in the console”](#).

Configuring Google Cloud

If you use **Default OAuth app** authentication, skip this section and see [the section called “Setting up the connector in Amazon Quick”](#).

For Custom OAuth app authentication, complete the following steps in the Google Cloud Console before you configure Amazon Quick. When you enable the API in step 3, search for and enable the **Google Docs API**.

Create an OAuth client in Google Cloud Console

Create an OAuth client in the Google Cloud Console to get the client credentials that you need for Amazon Quick. For more information, see [Using OAuth 2.0 to Access Google APIs](#) on the Google website.

1. Sign in to the [Google Cloud Console](#) on the Google website.
2. Create a new project or select an existing project.
3. In the left navigation pane, choose **APIs & Services**, then choose **Library**. Search for the API that your integration requires and choose **Enable**.
4. Choose **OAuth consent screen** and choose **Get Started**.
5. Configure the consent screen:
 - Enter an **App Name** and select a **User support email**.
 - For **Audience**, choose **Internal** (your organization only) or **External** (any Google user).
 - Add developer contact details and choose **Create**.
6. Choose **Create OAuth client**.
7. Configure the client:
 - For **Application type**, choose **Web application**.
 - Enter a **Name** for your client.
 - Under **Authorized redirect URIs**, add the Amazon Quick callback URL:
`https://{region}.quicksight.aws.amazon.com/sn/oauthcallback`
8. Choose **Create**.
9. Record the following values. You need them when you configure Amazon Quick.
 - **Client ID**
 - **Client secret**

Recommended scopes

The following scopes are requested when you connect to Google Docs.

Google Docs recommended scopes

Scope	Description
https://www.googleapis.com/auth/documents	Reads and writes Google Docs documents.
https://www.googleapis.com/auth/drive	Reads and writes files in Google Drive.
https://www.googleapis.com/auth/drive.readonly	Reads files in Google Drive.
openid	Authenticates the user's identity.
email	Reads the user's email address.

Note

The `openid` and `email` scopes are automatically included in the OAuth consent flow for user authentication. They are not required for Service-to-Service OAuth.

Setting up the connector in Amazon Quick

Connect from the Available tab

If you want to use Default OAuth app authentication, you can connect directly from the **Available** tab without additional configuration.

1. In the Amazon Quick console, choose **Connectors**.
2. On the **Available** tab, find **Google Docs** and choose **Connect**.
3. Complete the Google sign-in flow and grant the requested permissions.

To configure a connector with Custom OAuth app instead, use the **Create for your team** tab as described below.

Create from the Create for your team tab

After you complete any required Google Cloud configuration, create the connector in Amazon Quick.

1. In the Amazon Quick console, choose **Connectors**.
2. Choose the **Create for your team** tab.
3. Find and choose **Google Docs**.

Note

If a Google Docs connector already exists, a dialog appears with your existing connectors. To use an existing connector, choose it. To create a new one, choose **No, create new**.

4. Enter a **Name** for your connector. Optionally, choose **+ Add Description** to add a description.
5. For **Connection type**, choose **Public network**.
6. For **OAuth Configuration**, choose one of the following authentication methods and configure the required fields.
 - a. For **Default OAuth app**:

No additional credentials are needed. Choose **Next** to continue.
 - b. For **Custom OAuth app**, configure the following fields:
 - **Base URL** (Optional) – The Google Docs API base URL. For example: `https://docs.googleapis.com`
 - **Client ID** – The client ID from your Google Cloud OAuth client.
 - **Client secret** – The client secret from your Google Cloud OAuth client.
 - **Token URL** – The token endpoint. For example: `https://oauth2.googleapis.com/token`
 - **Authorization URL** – The authorization endpoint. For example: `https://accounts.google.com/o/oauth2/v2/auth`
 - **Redirect URL** – Pre-filled with the Amazon Quick callback URL.
7. Choose **Next**.

8. If you chose **Default OAuth app** or **Custom OAuth app**, a Google authorization window opens. Review the requested permissions and choose **Allow**.
9. On the **Review** page, review the available actions for the connector. Choose **Next**.
10. On the **Publish** page, choose who can access the connector. You can enable access for everyone in your organization or search for specific teams or groups.
11. Choose **Publish**.

Available actions

After you set up the connector, the following actions are available.

Google Docs available actions

Category	Action	Description
Documents	Get Document	Retrieves a document by ID.
Documents	Get Document Plaintext	Retrieves the plain text content of a document.
Documents	Search Documents	Searches for documents.
Documents	Create Document	Creates a new document.
Documents	Create Document Markdown	Creates a new document from Markdown content.
Documents	Copy Document	Creates a copy of an existing document.
Documents	Update Document	Updates an existing document.
Documents	Update Document Markdown	Updates a document with Markdown content.
Documents	Update Document Section	Updates a specific section of a document with Markdown.
Documents	Update Document Style	Updates the style of a document.
Documents	Export Pdf	Exports a document as a PDF file.

Category	Action	Description
Content	Insert Text	Inserts text at a specified location.
Content	Update Text	Replaces all occurrences of specified text.
Content	Delete Content Range	Deletes content within a specified range.
Content	Insert Inline Image	Inserts an inline image into a document.
Content	Create Table	Inserts a table into a document.
Content	Create Paragraph Bullets	Adds bullet formatting to paragraphs.
Content	Create Named Range	Creates a named range in a document.
Content	Create Page Break	Inserts a page break.
Structure	Create Header	Creates a header in a document.
Structure	Create Footer	Creates a footer in a document.

Note

The actions that you can use depend on the documents accessible to the authenticated user.

Managing and troubleshooting

To edit, share, or delete your connector, see [the section called “Managing existing integrations”](#).

Common Google authentication issues

- **Sign-in fails (Default OAuth app or Custom OAuth app)** – Verify that your Google account is active and that you can sign in to [the Google website](#) directly. For Custom OAuth app, confirm that the redirect URI in your Google Cloud OAuth client matches the Amazon Quick callback URL.
- **App blocked by administrator** – If your Google Workspace administrator restricts third-party app access, you might see an error when you attempt to sign in. Contact your Google Workspace administrator to allow the Amazon Quick app.

- **Invalid client credentials (Custom OAuth app)** – Verify that the Client ID and Client secret match the values in your Google Cloud OAuth client.
- **Authentication popup fails** – Verify that your browser allows popups from the Amazon Quick console domain. Try using a different browser or clearing your browser cache.
- **Permissions revoked** – If you previously revoked Amazon Quick access from your Google Account permissions settings, you need to re-authenticate by editing the connector and signing in again.
- **Google API rate limiting** – Google might limit requests during high usage periods. If actions fail, retry after a few minutes.

Google Docs-specific issues

- **Google Docs API not enabled** – Verify that the Google Docs API is enabled in your Google Cloud project under **APIs & Services, Library**.

Google Drive integration

With the Google Drive integration, you can manage files and folders through actions, and index your Google Drive content into knowledge bases for AI-powered search and Q&A.

What you can do

Connector

Upload and download files, manage folders, search content, and track changes through the Google Drive API.

Knowledge base

Index Google Drive document libraries, folders, and files. Amazon Quick agents can then search and answer questions about the content. Amazon Quick supports two setup methods: user-managed setup for quick configuration, and admin-managed setup with service credentials for document-level access control.

Google Drive action connector

With the Google Drive connector, you can access Google Drive directly in Amazon Quick through natural language. You can upload and download files, manage folders, search content, and track changes without leaving Amazon Quick.

Amazon Quick supports multiple authentication methods for the Google Drive connector. Choose the method that best fits your organization's security requirements.

- **Default OAuth app** – Uses an AWS-managed OAuth application. No additional credentials are needed. You authenticate directly with your Google account.
- **Custom OAuth app** – Uses a customer-managed OAuth client created in the Google Cloud Console. This option gives your organization full control over the OAuth configuration.

For more information about the authentication methods that Amazon Quick supports, see [the section called “Authentication methods”](#).

Before you begin

Make sure that you have the following before you set up the integration.

- A Google account with access to Google Drive.
- For **Custom OAuth app**: Access to the [Google Cloud Console](#) on the Google website with permissions to create OAuth clients.
- For Amazon Quick subscription requirements, see [the section called “Set up integrations in the console”](#).

Configuring Google Cloud

If you use **Default OAuth app** authentication, skip this section and see [the section called “Setting up the connector in Amazon Quick”](#).

For Custom OAuth app authentication, complete the following steps in the Google Cloud Console before you configure Amazon Quick. When you enable the API in step 3, search for and enable the **Google Drive API**.

Create an OAuth client in Google Cloud Console

Create an OAuth client in the Google Cloud Console to get the client credentials that you need for Amazon Quick. For more information, see [Using OAuth 2.0 to Access Google APIs](#) on the Google website.

1. Sign in to the [Google Cloud Console](#) on the Google website.
2. Create a new project or select an existing project.
3. In the left navigation pane, choose **APIs & Services**, then choose **Library**. Search for the API that your integration requires and choose **Enable**.
4. Choose **OAuth consent screen** and choose **Get Started**.
5. Configure the consent screen:
 - Enter an **App Name** and select a **User support email**.
 - For **Audience**, choose **Internal** (your organization only) or **External** (any Google user).
 - Add developer contact details and choose **Create**.
6. Choose **Create OAuth client**.
7. Configure the client:
 - For **Application type**, choose **Web application**.
 - Enter a **Name** for your client.
 - Under **Authorized redirect URIs**, add the Amazon Quick callback URL:
`https://{region}.quicksight.aws.amazon.com/sn/oauthcallback`
8. Choose **Create**.
9. Record the following values. You need them when you configure Amazon Quick.
 - **Client ID**
 - **Client secret**

Recommended scopes

The following scopes are requested when you connect to Google Drive.

Google Drive recommended scopes

Scope	Description
<code>https://www.googleapis.com/auth/drive</code>	Reads and writes files in Google Drive.
<code>https://www.googleapis.com/auth/drive.readonly</code>	Reads files in Google Drive.
<code>https://www.googleapis.com/auth/drive.metadata.readonly</code>	Reads file metadata in Google Drive.
<code>openid</code>	Authenticates the user's identity.
<code>email</code>	Reads the user's email address.

Note

The `openid` and `email` scopes are automatically included in the OAuth consent flow for user authentication. They are not required for Service-to-Service OAuth.

Setting up the connector in Amazon Quick

Connect from the Available tab

If you want to use Default OAuth app authentication, you can connect directly from the **Available** tab without additional configuration.

1. In the Amazon Quick console, choose **Connectors**.
2. On the **Available** tab, find **Google Drive** and choose **Connect**.
3. Complete the Google sign-in flow and grant the requested permissions.

To configure a connector with Custom OAuth app instead, use the **Create for your team** tab as described below.

Create from the Create for your team tab

After you complete any required Google Cloud configuration, create the connector in Amazon Quick.

1. In the Amazon Quick console, choose **Connectors**.
2. Choose the **Create for your team** tab.
3. Find and choose **Google Drive**.

Note

If a Google Drive connector already exists, a dialog appears with your existing connectors. To use an existing connector, choose it. To create a new one, choose **No, create new**.

4. Enter a **Name** for your connector. Optionally, choose **+ Add Description** to add a description.
5. For **Connection type**, choose **Public network**.
6. For **OAuth Configuration**, choose one of the following authentication methods and configure the required fields.
 - a. For **Default OAuth app**:

No additional credentials are needed. Choose **Next** to continue.
 - b. For **Custom OAuth app**, configure the following fields:
 - **Base URL** (Optional) – The Google Drive API base URL. For example: `https://www.googleapis.com/drive`
 - **Client ID** – The client ID from your Google Cloud OAuth client.
 - **Client secret** – The client secret from your Google Cloud OAuth client.
 - **Token URL** – The token endpoint. For example: `https://oauth2.googleapis.com/token`
 - **Authorization URL** – The authorization endpoint. For example: `https://accounts.google.com/o/oauth2/v2/auth`
 - **Redirect URL** – Pre-filled with the Amazon Quick callback URL.
7. Choose **Next**.

8. If you chose **Default OAuth app** or **Custom OAuth app**, a Google authorization window opens. Review the requested permissions and choose **Allow**.
9. On the **Review** page, review the available actions for the connector. Choose **Next**.
10. On the **Publish** page, choose who can access the connector. You can enable access for everyone in your organization or search for specific teams or groups.
11. Choose **Publish**.

Available actions

After you set up the connector, the following actions are available.

Google Drive available actions

Category	Action	Description
Files	Get File	Downloads a file from Google Drive.
Files	Search Files	Searches for files by name or content.
Files	Get File Metadata	Retrieves metadata for a file.
Files	Import File	Uploads a file to Google Drive.
Files	Upload File From Url	Uploads a file from a URL.
Files	Create File	Creates a new file.
Files	Create File From Text	Creates a file from text content.
Files	Put File	Replaces the content of an existing file.
Files	Update File Content	Edits the content of an existing file.
Files	Update File Location	Moves a file to a different folder.
Files	Copy File	Creates a copy of a file.
Files	Export File	Exports a Google Workspace file to a different format.

Category	Action	Description
Files	Delete File	Moves a file to the trash.
Folders	Create Folder	Creates a new folder.
Folders	Search Folders	Searches for folders by name.
Permissions	List Permissions	Lists permissions for a file or folder.
Comments	List Comments	Lists comments on a file.
Changes	List Changes	Lists recent changes to files.
Changes	Get Changes Start Page Token	Retrieves a start page token for tracking changes.
Drive	Get Drive Info	Retrieves information about the user's Drive.

Note

The actions that you can use depend on the files and folders accessible to the authenticated user.

Managing and troubleshooting

To edit, share, or delete your connector, see [the section called “Managing existing integrations”](#).

Common Google authentication issues

- **Sign-in fails (Default OAuth app or Custom OAuth app)** – Verify that your Google account is active and that you can sign in to [the Google website](#) directly. For Custom OAuth app, confirm that the redirect URI in your Google Cloud OAuth client matches the Amazon Quick callback URL.
- **App blocked by administrator** – If your Google Workspace administrator restricts third-party app access, you might see an error when you attempt to sign in. Contact your Google Workspace administrator to allow the Amazon Quick app.

- **Invalid client credentials (Custom OAuth app)** – Verify that the Client ID and Client secret match the values in your Google Cloud OAuth client.
- **Authentication popup fails** – Verify that your browser allows popups from the Amazon Quick console domain. Try using a different browser or clearing your browser cache.
- **Permissions revoked** – If you previously revoked Amazon Quick access from your Google Account permissions settings, you need to re-authenticate by editing the connector and signing in again.
- **Google API rate limiting** – Google might limit requests during high usage periods. If actions fail, retry after a few minutes.

Google Drive-specific issues

- **Google Drive API not enabled** – Verify that the Google Drive API is enabled in your Google Cloud project under **APIs & Services, Library**.

Google Drive knowledge base integration

With the Google Drive knowledge base integration, you can index your Google Drive content. Your Amazon Quick agents can then search this content and answer questions about it.

Amazon Quick supports two authentication methods for connecting to Google Drive:

- **User-managed setup** – You sign in to Google Drive directly to authorize the connection. This is the simplest way to get started. For more information, see [the section called “User-managed setup”](#).
- **Admin-managed setup (service credentials)** – A Google Workspace administrator creates a service account with domain-wide delegation to authorize the connection. A key benefit of admin-managed setup is built-in document-level access control (ACL). Amazon Quick automatically syncs access control lists from Google Drive. It verifies each user's permissions at query time, so users see answers only from documents that they are authorized to access. For more information, see [the section called “Admin-managed Google Drive knowledge base setup”](#).

After you connect, Amazon Quick indexes your Google Drive files and folders into a knowledge base. Your Amazon Quick agents can then search this content and generate answers that are grounded in your Google Drive data.

Prerequisites

Before you set up the Google Drive knowledge base integration, make sure that you have the following:

- A Google account with Google Drive access.
- For subscription requirements, see [the section called “Set up integrations in the console”](#).

For admin-managed setup, additional prerequisites apply. For more information, see [the section called “Admin-managed Google Drive knowledge base setup”](#).

Supported content types

The Google Drive connector supports all of the common file types that Amazon Quick knowledge bases support. These include PDF, Word, Excel, PowerPoint, and text files. The connector also supports the following Google Workspace-specific formats:

- Google Docs
- Google Sheets
- Google Slides

For more information about supported file types, size limits, and content processing options, see [Common configuration settings](#).

User-managed setup

With user-managed setup, you sign in to Google Drive directly to authorize the connection. Amazon Quick handles authentication through a managed OAuth flow. No Google Cloud project, service account, or domain-wide delegation is required.

Prerequisites

Before you set up a user-managed Google Drive knowledge base, verify the following:

- You have a Google account with access to Google Drive.
- Your Google Workspace administrator allows third-party app access, or can allow the Amazon Quick app on your behalf.
- Your browser allows popups from the Amazon Quick console domain.

Note

If your organization restricts third-party app access in Google Workspace, your Google Workspace administrator might need to allow the Amazon Quick app before users can sign in. Contact your Google Workspace administrator if you encounter an error during sign-in.

Note

User-managed setup does not support document-level access control (ACL). ACL is a mechanism that controls which users can access specific documents. If you need document-level access control, use [Admin-managed Google Drive knowledge base setup](#) instead.

Permissions granted during consent

When you authorize the connection, Amazon Quick requests the following permissions from your Google account:

See and download all your Google Drive files

- Allows Amazon Quick to see your Google Drive files
- Allows Amazon Quick to download your files
- Allows Amazon Quick to see the names and email addresses of people you share files with

See information about your Google Drive files

- Allows Amazon Quick to see the titles and descriptions of your files
- Allows Amazon Quick to see the names and email addresses of people you share files with
- Allows Amazon Quick to see your folders and how files are organized

Note

You can review and remove this access at any time from your Google Account permissions settings.

Set up a Google Drive knowledge base

To create a user-managed Google Drive knowledge base, complete the following steps in the Amazon Quick console.

1. In the Amazon Quick console, choose **Knowledge**.
2. Find **Google Drive** and choose the **Add (+)** icon.
3. In the **Create Google Drive knowledge base** dialog, under **Authentication method**, choose **Sign in to Google Drive** and complete the Google sign-in and consent flow.
4. Under **Create knowledge base**, enter a name and an optional description for your knowledge base.
5. In the **Content** section, choose **Add content** and select the Google Drive files and folders that you want to index. You can browse content from your personal drive, files shared with you, and shared drives in your organization.
6. Choose **Create**.

After you choose **Create**, the data sync starts automatically.

Access controls

Important

When Amazon Quick indexes Google Drive content through user-managed setup, it does not sync access control lists (ACLs) from Google Drive. All indexed content is accessible to any user who has access to the knowledge base in Amazon Quick, regardless of their permissions in Google Drive. Carefully review which content you include when you create a knowledge base.

If you require document-level access control, use the [Admin-managed Google Drive knowledge base setup](#) instead.

Manage and troubleshoot your integration

For instructions on editing, sharing, or deleting your integration, see [the section called “Managing existing integrations”](#).

For more information about knowledge base troubleshooting, including sync issues and missing documents, see [the section called “Troubleshooting knowledge bases”](#).

Google Drive-specific issues

- **App blocked by administrator** – If your Google Workspace administrator restricts third-party app access, you might see an error when you attempt to sign in. Contact your Google Workspace administrator to allow the Amazon Quick app.
- **Authentication popup fails** – Verify that your browser allows popups from the Amazon Quick console domain. Try using a different browser or clearing your browser cache.
- **Permissions revoked** – If you previously revoked Amazon Quick access from your Google Account permissions settings, you need to re-authenticate by editing the integration and signing in again.
- **Missing content** – Verify that the Google account that you used for authentication has access to the files and folders that you selected. Content that was shared with you after the initial sync requires a resync to be indexed.
- **Google API rate limiting** – Google Drive might limit requests during high usage periods. If syncs fail or are incomplete, retry during off-peak hours.

Known limitations

- Document-level access control (ACL) is not supported with user-managed setup. If you require document-level access control, use [Admin-managed Google Drive knowledge base setup](#).
- Synchronization of file comments is not supported.

Admin-managed Google Drive knowledge base setup

With admin-managed setup, a Google Workspace administrator creates a service account and delegates domain-wide access. Individual users don't need to authorize through sign-in.

Admin-managed setup includes built-in document-level access control list (ACL) support. Amazon Quick automatically syncs ACLs from Google Drive and verifies each user's permissions at query time.

For more information about ACL best practices, see [the section called “Best practices for managing ACLs in knowledge bases”](#).

Prerequisites

Make sure that you have the following before you set up the integration.

- Administrator access to your organization's Google Workspace.
- An Amazon Quick enterprise user account. Administrator access is not required.
- A Google Workspace account with an email domain that matches the email domain that is used for your Amazon Quick identity.
- For subscription requirements, see [the section called "Set up integrations in the console"](#).

Setup overview

The setup involves the following phases:

1. **Configure Google Workspace** – Create a Google Cloud service account with read-only API access and domain-wide delegation. Then create a dedicated admin user for the service account to impersonate. For more information, see [the section called "Configure Google Workspace"](#).
2. **Create the knowledge base in Amazon Quick** – Create a Google Drive knowledge base by using the service account credentials from Phase 1. For more information, see [the section called "Creating a knowledge base in Amazon Quick"](#).

Document-level access control is automatically enabled for all admin-managed knowledge bases. For more information about how access controls work, see [the section called "Document-level access controls"](#).

Configure Google Workspace

To connect Amazon Quick to Google Drive, complete the following tasks in the Google Cloud console and Google Workspace Admin Console. You create a Google Cloud project, turn on the required APIs, generate service account credentials, and configure domain-wide delegation. You also create a dedicated admin user for the service account to impersonate.

Prerequisites

Before you begin, make sure that you have the following:

- A Google Workspace account with administrator access

- Permission to create projects in the Google Cloud console

Creating a Google Cloud project

1. Open the Google Cloud console.
2. From the project selector at the top of the page, choose **New Project**.
3. Enter a project name, then choose **Create**.
4. After the project is created, choose **Select Project** to switch to it. This might take a few moments.

Turning on the required APIs

Amazon Quick requires three Google APIs. Turn on each one from the API Library.

1. In the navigation menu, choose **APIs & Services**, then choose **Library**.
2. Search for each of the following APIs and choose **Enable**:
 - Google Drive API
 - Google Drive Activity API
 - Admin SDK API

Creating the service account

1. In the navigation menu, choose **APIs & Services**, then choose **Credentials**.
2. Choose **Create Credentials**, then choose **Service account**.
3. Enter a name and optional description for the service account, then choose **Done**.

Generating a private key

1. On the **Credentials** page, choose the service account you created.
2. Choose the **Keys** tab, then choose **Add Key, Create new key**.
3. Confirm that **JSON** is selected, then choose **Create**.

The browser downloads a JSON file containing the private key. Store this file securely. You upload it to Amazon Quick in a later step.

 Note

If you receive an error stating that service account key creation is disabled by an organization policy, see [the section called “Resolving organization policy restrictions”](#).

Recording the service account unique ID

1. On the service account detail page, choose the **Details** tab.
2. Copy the value in the **Unique ID** field. You need this value when you configure domain-wide delegation.

Configuring domain-wide delegation

Domain-wide delegation allows the service account to access Google Workspace data on behalf of users in your organization.

1. On the service account detail page, expand **Advanced settings**.
2. Choose **View Google Workspace Admin Console**. The admin console opens in a new tab.
3. In the admin console navigation pane, choose **Security, Access and data control, API controls**.
4. Choose **Manage Domain Wide Delegation**, then choose **Add new**.
5. For **Client ID**, enter the unique ID you copied earlier.
6. For **OAuth scopes**, enter the following comma-separated values:

```
https://www.googleapis.com/auth/drive.readonly,https://www.googleapis.com/auth/drive.metadata.readonly,https://www.googleapis.com/auth/admin.directory.user.readonly,https://www.googleapis.com/auth/admin.directory.group.readonly,https://www.googleapis.com/auth/cloud-platform,https://www.googleapis.com/auth/forms.body.readonly
```

7. Choose **Authorize**.

Creating a delegated admin user

The service account acts on behalf of a Google Workspace admin user. Create a dedicated user for this purpose and assign the minimum required roles.

1. In the Google Workspace Admin Console, choose **Directory**, then choose **Users**.
2. Choose **Add new user**.
3. Enter a first name, last name, and primary email address for the new user, then choose **Add new user**.
4. Choose **Done**.
5. From the user list, choose the user you created. If the user does not appear, refresh the page.
6. On the user detail page, expand the **Admin roles and privileges** section.
7. Under **Roles**, assign the following roles:
 - Groups Reader
 - User Management Admin
 - Storage Admin
8. Choose **Save**.

Record the email address of this user. You need it when you create the knowledge base in Amazon Quick.

Troubleshooting the Google Workspace configuration

Resolving organization policy restrictions

If you receive the following error when creating a service account key:

The organization policy constraint `iam.disableServiceAccountKeyCreation` is enforced on your organization.

Note

For Google Cloud organizations created on or after May 3, 2024, this constraint is enforced by default.

You must override the policy for your project.

1. Open the Google Cloud console and confirm that the correct project is selected.
2. In the navigation menu, choose **IAM & Admin**, then choose **Organization Policies**.
3. In the **Filter** field, enter `iam.disableServiceAccountKeyCreation`. Then, in the policy list, choose **Disable service account key creation**.
4. Choose **Manage policy**.

 Note

If **Manage policy** is unavailable, you need the Organization Policy Administrator role (`roles/orgpolicy.policyAdmin`) at the organization level. See [the section called "Granting the Organization Policy Administrator role"](#).

5. In the **Policy source** section, ensure that **Override parent's policy** is selected.
6. Under **Enforcement**, turn off enforcement for this organization policy constraint.
7. Choose **Set policy**.

The change can take several minutes to propagate.

Granting the Organization Policy Administrator role

The Organization Policy Administrator role (`roles/orgpolicy.policyAdmin`) must be granted at the organization level, not the project level. It does not appear in the role list when assigning roles to a project.

To grant this role, select your organization (not a project) from the project selector in the Google Cloud console. Then, choose **IAM & Admin**, **IAM**, and assign the role to your account. For detailed instructions, see [Manage access to projects, folders, and organizations](#) in the Google Cloud documentation.

The role assignment can take several minutes to propagate.

Creating a knowledge base in Amazon Quick

In this phase, you create a knowledge base in Amazon Quick and provide the service account credentials from the Google Workspace configuration. Any enterprise user can complete this phase. Amazon Quick administrator access is not required.

If a Google Workspace administrator completed the Google Workspace configuration on your behalf, you need the JSON key file and the delegated admin email address before you proceed.

Setting up the knowledge base

1. In the Amazon Quick console, choose **Knowledge**.
2. Under **Knowledge bases**, find **Google Drive**, and then choose the **Add (+)** icon.
3. In the **Create Google Drive knowledge base** dialog, choose **Have admin credentials? Configure document-level access control**.
4. In the **Connected account** dropdown, choose **Add account**.
5. For **Name**, enter a name for the connection. Use a descriptive name such as your Google Workspace domain.

Important

You cannot change the connection name after you save it.

6. Choose **Upload .JSON key**, and then choose the JSON file that you downloaded during the Google Workspace configuration.
7. For **Google workspace admin email**, enter the email address of the delegated admin user that you created during the Google Workspace configuration.
8. Choose **Next**.

Choosing content to sync

1. Enter a **Name** and optional **Description** for your knowledge base.
2. Choose which Google Drive content to include:
 - **My Drive (all users)** – Includes files from all users' My Drive in your organization.
 - **Shared with me (all users)** – Includes files that are shared with your users.
 - **Shared drives** – All shared drives sync by default. To include or exclude specific drives, use the **Filter type** dropdown and **Add shared drive IDs** field. You can enter 1 to 100 shared drive IDs.
3. Choose **Next** to configure advanced settings.

Configuring advanced settings

In the **Advanced settings** step, you can configure optional settings for the knowledge base.

Filter content by date

Limit which documents are crawled based on their last modified date. The start date defaults to one year before today. You can change or clear the start date, and optionally set an end date.

Multi-media content, file size, and file patterns

Choose which content types to include in the knowledge base.

- **Visual content in documents** – Extracts and indexes visual elements from supported document formats. This option is enabled by default.
- **Audio files** – Transcribes and indexes audio files.
- **Video files** – Transcribes and indexes video files.

Choose **Create** to create the knowledge base. After you choose **Create**, the data sync starts automatically.

Managing and troubleshooting

To edit, share, or delete your integration, see [the section called “Managing existing integrations”](#).

For information about knowledge base troubleshooting, including sync issues and missing documents, see [the section called “Troubleshooting knowledge bases”](#).

For information about sync reports, sync schedules, and verifying document-level access controls, see [the section called “Sync reports and observability”](#).

Admin-managed setup issues

- **Google API rate limiting** – Google Drive might throttle requests during high usage periods. If syncs fail or are incomplete, retry during off-peak hours.
- **SSL certificate errors** – If you receive an error about SSL certificate errors when you create your knowledge base, verify the OAuth scopes that you configured during domain-wide delegation.

Document-level access controls

Admin-managed Google Drive knowledge bases include built-in document-level access control. Amazon Quick syncs access control lists (ACLs) from Google Drive during each crawl and verifies

each user's permissions at query time, so users only see answers from documents that they are authorized to access.

How it works

When a user submits a query to an Amazon Quick agent that uses an admin-managed Google Drive knowledge base, the system enforces access controls in two stages:

1. **Pre-retrieval filtering** – Amazon Quick performs a semantic search against the vector index to find the most relevant document passages. The system applies access control lists that are already stored in the index. This produces a preliminary set of candidate documents. This stage is necessary because real-time API calls for every document in the index would be too costly at scale.
2. **Real-time verification** – The system verifies the candidate documents in real time by calling the Google Drive APIs. It uses the service account credential that the administrator provided to generate user-specific access tokens through impersonation. Google Drive maintains the source of truth for access control lists that are associated with each document. The system removes any documents that the user is not authorized to access from the result set.

The system passes only the verified and authorized document passages to the model as context. The model uses this knowledge to generate a response. This two-stage approach provides document-level access control guarantees and maintains performance at scale.

Enable ACL management

Document-level access control is automatically enabled for all admin-managed knowledge bases. No additional configuration is required.

For more information about ACL best practices, see [the section called “Best practices for managing ACLs in knowledge bases”](#).

Known limitations

- File comments synchronization is not supported.

For more information about general ACL limitations and best practices, see [the section called “Best practices for managing ACLs in knowledge bases”](#).

Google Meet integration

With the Google Meet connector, you can access Google Meet directly in Amazon Quick through natural language. You can schedule meetings, view conference records, retrieve transcripts and recordings, and manage meeting spaces without leaving Amazon Quick.

Amazon Quick supports multiple authentication methods for Google Meet. Choose the method that best fits your organization's security requirements.

- **Default OAuth app** – Uses an AWS-managed OAuth application. No additional credentials are needed. You authenticate directly with your Google account.
- **Custom OAuth app** – Uses a customer-managed OAuth client created in the Google Cloud Console. This option gives your organization full control over the OAuth configuration.

For more information about the authentication methods that Amazon Quick supports, see [the section called “Authentication methods”](#).

Before you begin

Make sure that you have the following before you set up the integration.

- A Google account with access to Google Meet.
- For **Custom OAuth app**: Access to the [Google Cloud Console](#) on the Google website with permissions to create OAuth clients.
- For Amazon Quick subscription requirements, see [the section called “Set up integrations in the console”](#).

Configuring Google Cloud

If you use **Default OAuth app** authentication, skip this section and see [the section called “Setting up the connector in Amazon Quick”](#).

For Custom OAuth app authentication, complete the following steps in the Google Cloud Console before you configure Amazon Quick. When you enable the API in step 3, search for and enable the **Google Meet REST API**.

Create an OAuth client in Google Cloud Console

Create an OAuth client in the Google Cloud Console to get the client credentials that you need for Amazon Quick. For more information, see [Using OAuth 2.0 to Access Google APIs](#) on the Google website.

1. Sign in to the [Google Cloud Console](#) on the Google website.
2. Create a new project or select an existing project.
3. In the left navigation pane, choose **APIs & Services**, then choose **Library**. Search for the API that your integration requires and choose **Enable**.
4. Choose **OAuth consent screen** and choose **Get Started**.
5. Configure the consent screen:
 - Enter an **App Name** and select a **User support email**.
 - For **Audience**, choose **Internal** (your organization only) or **External** (any Google user).
 - Add developer contact details and choose **Create**.
6. Choose **Create OAuth client**.
7. Configure the client:
 - For **Application type**, choose **Web application**.
 - Enter a **Name** for your client.
 - Under **Authorized redirect URIs**, add the Amazon Quick callback URL:
`https://{region}.quicksight.aws.amazon.com/sn/oauthcallback`
8. Choose **Create**.
9. Record the following values. You need them when you configure Amazon Quick.
 - **Client ID**
 - **Client secret**

Recommended scopes

The following scopes are requested when you connect to Google Meet.

Google Meet recommended scopes

Scope	Description
<code>https://www.googleapis.com/auth/meetings.space.created</code>	Creates meeting spaces.
<code>https://www.googleapis.com/auth/meetings.space.readonly</code>	Reads meeting space information.
<code>https://www.googleapis.com/auth/meetings.space.settings</code>	Manages meeting space settings.
<code>openid</code>	Authenticates the user's identity.
<code>email</code>	Reads the user's email address.

Note

The `openid` and `email` scopes are automatically included in the OAuth consent flow for user authentication. They are not required for Service-to-Service OAuth.

Setting up the connector in Amazon Quick

Connect from the Available tab

If you want to use Default OAuth app authentication, you can connect directly from the **Available** tab without additional configuration.

1. In the Amazon Quick console, choose **Connectors**.
2. On the **Available** tab, find **Google Meet** and choose **Connect**.
3. Complete the Google sign-in flow and grant the requested permissions.

To configure a connector with Custom OAuth app instead, use the **Create for your team** tab as described below.

Create from the Create for your team tab

After you complete any required Google Cloud configuration, create the connector in Amazon Quick.

1. In the Amazon Quick console, choose **Connectors**.
2. Choose the **Create for your team** tab.
3. Find and choose **Google Meet**.

Note

If a Google Meet connector already exists, a dialog appears with your existing connectors. To use an existing connector, choose it. To create a new one, choose **No, create new**.

4. Enter a **Name** for your connector. Optionally, choose **+ Add Description** to add a description.
5. For **Connection type**, choose **Public network**.
6. For **OAuth Configuration**, choose one of the following authentication methods and configure the required fields.
 - a. For **Default OAuth app**:

No additional credentials are needed. Choose **Next** to continue.
 - b. For **Custom OAuth app**, configure the following fields:
 - **Base URL** (Optional) – The Google Meet API base URL. For example: `https://meet.googleapis.com`
 - **Client ID** – The client ID from your Google Cloud OAuth client.
 - **Client secret** – The client secret from your Google Cloud OAuth client.
 - **Token URL** – The token endpoint. For example: `https://oauth2.googleapis.com/token`
 - **Authorization URL** – The authorization endpoint. For example: `https://accounts.google.com/o/oauth2/v2/auth`
 - **Redirect URL** – Pre-filled with the Amazon Quick callback URL.
7. Choose **Next**.

8. If you chose **Default OAuth app** or **Custom OAuth app**, a Google authorization window opens. Review the requested permissions and choose **Allow**.
9. On the **Review** page, review the available actions for the connector. Choose **Next**.
10. On the **Publish** page, choose who can access the connector. You can enable access for everyone in your organization or search for specific teams or groups.
11. Choose **Publish**.

Available actions

After you set up the connector, the following actions are available.

Google Meet available actions

Category	Action	Description
Meetings	Create Meeting	Creates a new meeting.
Meetings	Get Meeting	Retrieves details for a specific meeting.
Meetings	List Conference Records	Lists conference records for past meetings.
Spaces	Create Space	Creates a new meeting space.
Spaces	Get Space	Retrieves details for a meeting space.
Spaces	Update Space	Updates meeting space settings.
Participants	List Participants	Lists participants for a conference record.
Participants	List Participant Sessions	Lists sessions for a participant.
Participants	Get Participant Session	Retrieves details for a participant session.
Recordings	List Recordings	Lists recordings for a conference record.
Recordings	Get Recordings	Retrieves recordings by conference record ID.

Category	Action	Description
Transcripts	Get Transcripts	Retrieves transcripts by conference record ID.
Transcripts	List Transcript Entries	Lists individual entries in a transcript.
Transcripts	Get Transcript Entry	Retrieves a specific transcript entry.

Note

The actions that you can use depend on the meetings and conference records accessible to the authenticated user.

Managing and troubleshooting

To edit, share, or delete your connector, see [the section called “Managing existing integrations”](#).

Common Google authentication issues

- **Sign-in fails (Default OAuth app or Custom OAuth app)** – Verify that your Google account is active and that you can sign in to [the Google website](#) directly. For Custom OAuth app, confirm that the redirect URI in your Google Cloud OAuth client matches the Amazon Quick callback URL.
- **App blocked by administrator** – If your Google Workspace administrator restricts third-party app access, you might see an error when you attempt to sign in. Contact your Google Workspace administrator to allow the Amazon Quick app.
- **Invalid client credentials (Custom OAuth app)** – Verify that the Client ID and Client secret match the values in your Google Cloud OAuth client.
- **Authentication popup fails** – Verify that your browser allows popups from the Amazon Quick console domain. Try using a different browser or clearing your browser cache.
- **Permissions revoked** – If you previously revoked Amazon Quick access from your Google Account permissions settings, you need to re-authenticate by editing the connector and signing in again.
- **Google API rate limiting** – Google might limit requests during high usage periods. If actions fail, retry after a few minutes.

Google Meet-specific issues

- **Google Meet REST API not enabled** – Verify that the Google Meet REST API is enabled in your Google Cloud project under **APIs & Services, Library**.

Google Sheets integration

With the Google Sheets connector, you can access Google Sheets directly in Amazon Quick through natural language. You can create and edit spreadsheets, read and write cell ranges, format cells, and search for data without leaving Amazon Quick.

Amazon Quick supports multiple authentication methods for Google Sheets. Choose the method that best fits your organization's security requirements.

- **Default OAuth app** – Uses an AWS-managed OAuth application. No additional credentials are needed. You authenticate directly with your Google account.
- **Custom OAuth app** – Uses a customer-managed OAuth client created in the Google Cloud Console. This option gives your organization full control over the OAuth configuration.

For more information about the authentication methods that Amazon Quick supports, see [the section called “Authentication methods”](#).

Before you begin

Make sure that you have the following before you set up the integration.

- A Google account with access to Google Sheets.
- For **Custom OAuth app**: Access to the [Google Cloud Console](#) on the Google website with permissions to create OAuth clients.
- For Amazon Quick subscription requirements, see [the section called “Set up integrations in the console”](#).

Configuring Google Cloud

If you use **Default OAuth app** authentication, skip this section and see [the section called “Setting up the connector in Amazon Quick”](#).

For Custom OAuth app authentication, complete the following steps in the Google Cloud Console before you configure Amazon Quick. When you enable the API in step 3, search for and enable the **Google Sheets API**.

Create an OAuth client in Google Cloud Console

Create an OAuth client in the Google Cloud Console to get the client credentials that you need for Amazon Quick. For more information, see [Using OAuth 2.0 to Access Google APIs](#) on the Google website.

1. Sign in to the [Google Cloud Console](#) on the Google website.
2. Create a new project or select an existing project.
3. In the left navigation pane, choose **APIs & Services**, then choose **Library**. Search for the API that your integration requires and choose **Enable**.
4. Choose **OAuth consent screen** and choose **Get Started**.
5. Configure the consent screen:
 - Enter an **App Name** and select a **User support email**.
 - For **Audience**, choose **Internal** (your organization only) or **External** (any Google user).
 - Add developer contact details and choose **Create**.
6. Choose **Create OAuth client**.
7. Configure the client:
 - For **Application type**, choose **Web application**.
 - Enter a **Name** for your client.
 - Under **Authorized redirect URIs**, add the Amazon Quick callback URL:
`https://{region}.quicksight.aws.amazon.com/sn/oauthcallback`
8. Choose **Create**.
9. Record the following values. You need them when you configure Amazon Quick.
 - **Client ID**
 - **Client secret**

Recommended scopes

The following scopes are requested when you connect to Google Sheets.

Google Sheets recommended scopes

Scope	Description
https://www.googleapis.com/auth/spreadsheets	Reads and writes Google Sheets spreadsheets.
https://www.googleapis.com/auth/drive	Reads and writes files in Google Drive.
https://www.googleapis.com/auth/drive.readonly	Reads files in Google Drive.
openid	Authenticates the user's identity.
email	Reads the user's email address.

Note

The `openid` and `email` scopes are automatically included in the OAuth consent flow for user authentication. They are not required for Service-to-Service OAuth.

Setting up the connector in Amazon Quick

Connect from the Available tab

If you want to use Default OAuth app authentication, you can connect directly from the **Available** tab without additional configuration.

1. In the Amazon Quick console, choose **Connectors**.
2. On the **Available** tab, find **Google Sheets** and choose **Connect**.
3. Complete the Google sign-in flow and grant the requested permissions.

To configure a connector with Custom OAuth app instead, use the **Create for your team** tab as described below.

Create from the Create for your team tab

After you complete any required Google Cloud configuration, create the connector in Amazon Quick.

1. In the Amazon Quick console, choose **Connectors**.
2. Choose the **Create for your team** tab.
3. Find and choose **Google Sheets**.

Note

If a Google Sheets connector already exists, a dialog appears with your existing connectors. To use an existing connector, choose it. To create a new one, choose **No, create new**.

4. Enter a **Name** for your connector. Optionally, choose **+ Add Description** to add a description.
5. For **Connection type**, choose **Public network**.
6. For **OAuth Configuration**, choose one of the following authentication methods and configure the required fields.

a. For **Default OAuth app**:

No additional credentials are needed. Choose **Next** to continue.

b. For **Custom OAuth app**, configure the following fields:

- **Base URL** (Optional) – The Google Sheets API base URL. For example: `https://sheets.googleapis.com`
- **Client ID** – The client ID from your Google Cloud OAuth client.
- **Client secret** – The client secret from your Google Cloud OAuth client.
- **Token URL** – The token endpoint. For example: `https://oauth2.googleapis.com/token`
- **Authorization URL** – The authorization endpoint. For example: `https://accounts.google.com/o/oauth2/v2/auth`
- **Redirect URL** – Pre-filled with the Amazon Quick callback URL.

7. Choose **Next**.

8. If you chose **Default OAuth app** or **Custom OAuth app**, a Google authorization window opens. Review the requested permissions and choose **Allow**.
9. On the **Review** page, review the available actions for the connector. Choose **Next**.
10. On the **Publish** page, choose who can access the connector. You can enable access for everyone in your organization or search for specific teams or groups.
11. Choose **Publish**.

Available actions

After you set up the connector, the following actions are available.

Google Sheets available actions

Category	Action	Description
Values	Get Values	Reads values from a cell range.
Values	Batch Get Value	Reads values from multiple cell ranges in a single request.
Values	Get Values By Filter	Reads values using a data filter.
Values	Update Values	Writes values to a cell range.
Values	Batch Update Value	Writes values to multiple cell ranges in a single request.
Values	Append Values	Appends values after the last row of data in a range.
Values	Delete Values	Clears values from a cell range.
Values	Update Text	Finds and replaces text across a spreadsheet.
Rows	Create Row	Creates a new row in a spreadsheet.
Rows	Put Rows	Upserts rows in a spreadsheet.
Rows	Get Spreadsheet Row	Looks up a row by matching criteria.

Category	Action	Description
Spreadsheets	Create Spreadsheet	Creates a new spreadsheet.
Spreadsheets	Get Spreadsheet Info	Retrieves metadata for a spreadsheet.
Spreadsheets	Search Spreadsheets	Searches for spreadsheets.
Sheets	Create Sheet	Adds a new sheet to a spreadsheet.
Sheets	List Sheet Names	Lists sheet names in a spreadsheet.
Sheets	Update Sheet Properties	Updates sheet properties such as name or visibility.
Formatting	Update Cell Format	Formats cells in a range.
Formatting	Update Conditional Format Rules	Creates or updates conditional formatting rules.

Note

The actions that you can use depend on the spreadsheets accessible to the authenticated user.

Managing and troubleshooting

To edit, share, or delete your connector, see [the section called “Managing existing integrations”](#).

Common Google authentication issues

- **Sign-in fails (Default OAuth app or Custom OAuth app)** – Verify that your Google account is active and that you can sign in to [the Google website](#) directly. For Custom OAuth app, confirm that the redirect URI in your Google Cloud OAuth client matches the Amazon Quick callback URL.
- **App blocked by administrator** – If your Google Workspace administrator restricts third-party app access, you might see an error when you attempt to sign in. Contact your Google Workspace administrator to allow the Amazon Quick app.

- **Invalid client credentials (Custom OAuth app)** – Verify that the Client ID and Client secret match the values in your Google Cloud OAuth client.
- **Authentication popup fails** – Verify that your browser allows popups from the Amazon Quick console domain. Try using a different browser or clearing your browser cache.
- **Permissions revoked** – If you previously revoked Amazon Quick access from your Google Account permissions settings, you need to re-authenticate by editing the connector and signing in again.
- **Google API rate limiting** – Google might limit requests during high usage periods. If actions fail, retry after a few minutes.

Google Sheets-specific issues

- **Google Sheets API not enabled** – Verify that the Google Sheets API is enabled in your Google Cloud project under **APIs & Services, Library**.

Google Slides integration

With the Google Slides connector, you can access Google Slides directly in Amazon Quick through natural language. You can create presentations, edit slides, copy templates, and retrieve presentation content without leaving Amazon Quick.

Amazon Quick supports multiple authentication methods for Google Slides. Choose the method that best fits your organization's security requirements.

- **Default OAuth app** – Uses an AWS-managed OAuth application. No additional credentials are needed. You authenticate directly with your Google account.
- **Custom OAuth app** – Uses a customer-managed OAuth client created in the Google Cloud Console. This option gives your organization full control over the OAuth configuration.

For more information about the authentication methods that Amazon Quick supports, see [the section called “Authentication methods”](#).

Before you begin

Make sure that you have the following before you set up the integration.

- A Google account with access to Google Slides.

- For **Custom OAuth app**: Access to the [Google Cloud Console](#) on the Google website with permissions to create OAuth clients.
- For Amazon Quick subscription requirements, see [the section called "Set up integrations in the console"](#).

Configuring Google Cloud

If you use **Default OAuth app** authentication, skip this section and see [the section called "Setting up the connector in Amazon Quick"](#).

For Custom OAuth app authentication, complete the following steps in the Google Cloud Console before you configure Amazon Quick. When you enable the API in step 3, search for and enable the **Google Slides API**.

Create an OAuth client in Google Cloud Console

Create an OAuth client in the Google Cloud Console to get the client credentials that you need for Amazon Quick. For more information, see [Using OAuth 2.0 to Access Google APIs](#) on the Google website.

1. Sign in to the [Google Cloud Console](#) on the Google website.
2. Create a new project or select an existing project.
3. In the left navigation pane, choose **APIs & Services**, then choose **Library**. Search for the API that your integration requires and choose **Enable**.
4. Choose **OAuth consent screen** and choose **Get Started**.
5. Configure the consent screen:
 - Enter an **App Name** and select a **User support email**.
 - For **Audience**, choose **Internal** (your organization only) or **External** (any Google user).
 - Add developer contact details and choose **Create**.
6. Choose **Create OAuth client**.
7. Configure the client:
 - For **Application type**, choose **Web application**.
 - Enter a **Name** for your client.
 - Under **Authorized redirect URIs**, add the Amazon Quick callback URL:
`https://{region}.quicksight.aws.amazon.com/sn/oauthcallback`

8. Choose **Create**.
9. Record the following values. You need them when you configure Amazon Quick.
 - **Client ID**
 - **Client secret**

Recommended scopes

The following scopes are requested when you connect to Google Slides.

Google Slides recommended scopes

Scope	Description
<code>https://www.googleapis.com/auth/presentations</code>	Reads and writes Google Slides presentations.
<code>https://www.googleapis.com/auth/drive</code>	Reads and writes files in Google Drive.
<code>https://www.googleapis.com/auth/drive.readonly</code>	Reads files in Google Drive.
<code>openid</code>	Authenticates the user's identity.
<code>email</code>	Reads the user's email address.

Note

The `openid` and `email` scopes are automatically included in the OAuth consent flow for user authentication. They are not required for Service-to-Service OAuth.

Setting up the connector in Amazon Quick

Connect from the Available tab

If you want to use Default OAuth app authentication, you can connect directly from the **Available** tab without additional configuration.

1. In the Amazon Quick console, choose **Connectors**.
2. On the **Available** tab, find **Google Slides** and choose **Connect**.
3. Complete the Google sign-in flow and grant the requested permissions.

To configure a connector with Custom OAuth app instead, use the **Create for your team** tab as described below.

Create from the Create for your team tab

After you complete any required Google Cloud configuration, create the connector in Amazon Quick.

1. In the Amazon Quick console, choose **Connectors**.
2. Choose the **Create for your team** tab.
3. Find and choose **Google Slides**.

Note

If a Google Slides connector already exists, a dialog appears with your existing connectors. To use an existing connector, choose it. To create a new one, choose **No, create new**.

4. Enter a **Name** for your connector. Optionally, choose **+ Add Description** to add a description.
5. For **Connection type**, choose **Public network**.
6. For **OAuth Configuration**, choose one of the following authentication methods and configure the required fields.
 - a. For **Default OAuth app**:

No additional credentials are needed. Choose **Next** to continue.
 - b. For **Custom OAuth app**, configure the following fields:
 - **Base URL** (Optional) – The Google Slides API base URL. For example: `https://slides.googleapis.com`
 - **Client ID** – The client ID from your Google Cloud OAuth client.
 - **Client secret** – The client secret from your Google Cloud OAuth client.

- **Token URL** – The token endpoint. For example: `https://oauth2.googleapis.com/token`
- **Authorization URL** – The authorization endpoint. For example: `https://accounts.google.com/o/oauth2/v2/auth`
- **Redirect URL** – Pre-filled with the Amazon Quick callback URL.

7. Choose **Next**.
8. If you chose **Default OAuth app** or **Custom OAuth app**, a Google authorization window opens. Review the requested permissions and choose **Allow**.
9. On the **Review** page, review the available actions for the connector. Choose **Next**.
10. On the **Publish** page, choose who can access the connector. You can enable access for everyone in your organization or search for specific teams or groups.
11. Choose **Publish**.

Available actions

After you set up the connector, the following actions are available.

Google Slides available actions

Category	Action	Description
Presentations	Get Presentation	Retrieves a presentation by ID.
Presentations	Create Presentation	Creates a new presentation.
Presentations	Create Slides Markdown	Creates slides from Markdown content.
Presentations	Copy Presentation Template	Creates a presentation from a template.
Presentations	Batch Update Presentation	Applies multiple updates to a presentation in a single request.
Pages	Get Presentation Page	Retrieves a specific page from a presentation.

Note

The actions that you can use depend on the presentations accessible to the authenticated user.

Managing and troubleshooting

To edit, share, or delete your connector, see [the section called “Managing existing integrations”](#).

Common Google authentication issues

- **Sign-in fails (Default OAuth app or Custom OAuth app)** – Verify that your Google account is active and that you can sign in to [the Google website](#) directly. For Custom OAuth app, confirm that the redirect URI in your Google Cloud OAuth client matches the Amazon Quick callback URL.
- **App blocked by administrator** – If your Google Workspace administrator restricts third-party app access, you might see an error when you attempt to sign in. Contact your Google Workspace administrator to allow the Amazon Quick app.
- **Invalid client credentials (Custom OAuth app)** – Verify that the Client ID and Client secret match the values in your Google Cloud OAuth client.
- **Authentication popup fails** – Verify that your browser allows popups from the Amazon Quick console domain. Try using a different browser or clearing your browser cache.
- **Permissions revoked** – If you previously revoked Amazon Quick access from your Google Account permissions settings, you need to re-authenticate by editing the connector and signing in again.
- **Google API rate limiting** – Google might limit requests during high usage periods. If actions fail, retry after a few minutes.

Google Slides-specific issues

- **Google Slides API not enabled** – Verify that the Google Slides API is enabled in your Google Cloud project under **APIs & Services, Library**.

HG Insights integration

With the HG Insights connector, you can access technographic data, query company technology profiles, and retrieve market intelligence directly in Amazon Quick through natural language.

Amazon Quick supports multiple authentication methods for HG Insights. Choose the method that best fits your organization's security requirements.

- **Default OAuth app** – Uses an AWS-managed OAuth application. No additional credentials are needed. Users authenticate directly with their HG Insights account.
- **Custom OAuth app** – Uses a customer-managed OAuth application. This option gives your organization full control over the OAuth configuration.
- **Service-to-Service OAuth** – Uses client credentials for server-to-server authentication without user interaction. Suitable for automated workflows.
- **API Key** – Uses an API key for authentication.

For more information about the authentication methods that Amazon Quick supports, see [the section called “Authentication methods”](#).

Before you begin

Make sure that you have the following before you set up the integration.

- An active HG Insights account with access to the data products you want to query.
- For **Custom OAuth app** or **Service-to-Service OAuth**: OAuth credentials from your HG Insights account.
- For **API Key**: An HG Insights-issued API key.
- For Amazon Quick subscription requirements, see [the section called “Set up integrations in the console”](#).

Configuring HG Insights

If you are using **Default OAuth app** authentication, skip this section and proceed to [the section called “Setting up the connector in Amazon Quick”](#).

For Custom OAuth app or Service-to-Service OAuth authentication, register an OAuth client in your HG Insights Phoenix account and add the Amazon Quick callback URL `https://{region}.quicksight.aws.amazon.com/sn/oauthcallback` as a redirect URI. Replace `{region}` with your AWS Region (for example, `us-east-1`). For step-by-step instructions, see [Phoenix by HG Insights](#) documentation. Record the Client ID and Client Secret — you need them when you configure Amazon Quick. For API Key authentication, generate an API key from your Phoenix account.

Setting up the connector in Amazon Quick

Connect from the Available tab

If you want to use Default OAuth app authentication, you can connect directly from the **Available** tab without additional configuration.

1. In the Amazon Quick console, choose **Connectors**.
2. On the **Available** tab, find **HG Insights** and choose **Connect**.
3. Complete the HG Insights sign-in flow and grant the requested permissions.

To configure a connector with one of the other authentication methods, use the **Create for your team** tab as described below.

Create from the Create for your team tab

1. In the Amazon Quick console, choose **Connectors**.
2. Choose the **Create for your team** tab.
3. Find and choose **HG Insights**.
4. Enter a **Name** for your connector. Optionally, choose **+ Add Description** to add a description.
5. For **Connection type**, choose **Public network**.
6. For **OAuth Configuration**, choose one of the following authentication methods and configure the required fields.
 - a. For **Default OAuth app**:

No additional credentials are needed. Choose **Next** to continue.
 - b. For **Custom OAuth app**, configure the following fields:
 - **Client ID** – The client ID from your HG Insights OAuth app.
 - **Public OAuth client** (Optional) – Select this option if your HG Insights OAuth app is configured as a public client (no client secret).
 - **Client secret** – The client secret from your HG Insights OAuth app.
 - **Token URL** – The token endpoint. Default: `https://phoenix.hginsights.com/api/ai/token`
 - **Authorization URL** – The authorization endpoint. Default: `https://phoenix.hginsights.com/api/ai/authorize`

- **Redirect URL** – Pre-filled with the Amazon Quick callback URL.
- c. For **Service-to-Service OAuth**, configure the following fields:
 - **Client ID** – The client ID from your HG Insights OAuth app.
 - **Client secret** – The client secret from your HG Insights OAuth app.
 - **Token URL** – The token endpoint. Default: `https://phoenix.hginsights.com/api/ai/token`
- d. For **API Key**, configure the following fields:
 - **API Key** – The HG Insights API key.
 - **Email (Optional)** – The email address associated with the API key.
- 7. Choose **Next**.
- 8. If you chose **Default OAuth app** or **Custom OAuth app**, an HG Insights authorization window opens. Review the requested permissions and choose **Allow**.
- 9. On the **Review** page, review the available actions for the connector. Choose **Next**.
- 10. On the **Publish** page, choose who can access the connector. You can enable access for everyone in your organization or search for specific teams or groups.
- 11. Choose **Publish**.

Available actions

After you set up the connector, the actions exposed by HG Insights are available. To see the current set of actions for your connector, go to the connector's **Available actions** view in the Amazon Quick console.

Managing and troubleshooting

To edit, share, or delete your connector, see [the section called “Managing existing integrations”](#).

Authentication issues

- **Sign-in fails (Default OAuth app or Custom OAuth app)** – Verify that your HG Insights account is active. For Custom OAuth app, confirm that the redirect URI in your HG Insights OAuth app matches the Amazon Quick callback URL.
- **Invalid client credentials (Custom OAuth app or Service-to-Service OAuth)** – Verify that the Client ID and Client secret match the values in your HG Insights OAuth app.

- **API Key authentication fails** – Verify that the API key has not been revoked.

HubSpot integration

Access your HubSpot CRM data directly from Amazon Quick using natural language. With the HubSpot connector, you can retrieve and list contacts, companies, deals, and tickets without leaving Amazon Quick.

HubSpot is available as a built-in connector in Amazon Quick. To set up this integration, complete the following two steps. First, prepare your HubSpot account with the required access. Then, create the connector in Amazon Quick and authenticate with your HubSpot credentials.

This connector supports Custom OAuth app authentication. For information about the authentication methods that Amazon Quick supports, see [the section called "Authentication methods"](#).

Prerequisites

Before you set up the integration, verify that you have the following resources and access.

- An active HubSpot account with an admin role.
- Permissions to access CRM data including contacts, companies, deals, and tickets in your HubSpot account.
- A subscription that meets the requirements described in [the section called "Set up integrations in the console"](#).

Verifying your HubSpot account

Before you configure Amazon Quick, verify that your HubSpot account meets the following requirements.

- Your HubSpot account is active and you can sign in to the [HubSpot sign-in page](#) on the HubSpot website.
- The CRM data that you want to access, including contacts, companies, deals, and tickets, is accessible to your account.

Configuring your HubSpot account

The Custom OAuth app authentication method in Amazon Quick requires a MCP Auth app configured in your HubSpot developer settings.

To configure a MCP Auth App

1. Sign in to your HubSpot account.
2. Navigate to **Development**, and then **MCP Auth Apps**.
3. Choose **Create MCP auth app**.
4. On the **Create MCP Auth App** window, complete the following fields:
 - **Name** – Enter a name for your app.
 - **Description** – Enter a short description.
 - **Redirect URL** – Enter `https://region.quicksight.aws.amazon.com/sn/oauthcallback`. Replace *region* with your AWS Region, for example, `us-east-1`.
5. Choose **Create**. HubSpot creates the app and generates a client ID and client secret.

Important

Record the client secret and client ID.

The client ID and client secret serve as the Client ID and Client Secret when you configure the connector in Amazon Quick with the Custom OAuth app authentication method.

For more information, see [HubSpot MCP server](#) on the HubSpot website.

Setting up the integration in Amazon Quick

After you verify your HubSpot account access, create the connector in Amazon Quick.

1. In the Amazon Quick console, choose **Connectors**.
2. Choose the **Create for your team** tab.
3. Choose **HubSpot** from the options available.
4. In the **HubSpot** dialog, enter the following fields:
 - **Name** – A descriptive name for your HubSpot connector.

- **Description (Optional)** – Choose **+ Add Description** to add notes about how this connector will be used.
 - **Connection type** – Choose **Public network**.
5. For the **Custom OAuth app** authentication method, enter the following fields:
- **Base URL** – This field is auto-populated as `https://mcp.hubspot.com`.
 - **Client ID** – Enter the client ID from your HubSpot MCP Auth app.
 - **Client Secret** – Enter the client secret from your HubSpot MCP Auth app.
 - **Token URL** – This field is auto-populated as `https://api.hubapi.com/oauth/v1/token`.
 - **Authorization URL** – This field is auto-populated as `https://app.hubspot.com/oauth/authorize`.
 - **Redirect URL** – This field is auto-populated as `https://us-east-1.quicksight.aws.amazon.com/sn/oauthcallback`.
6. Choose **Next**, review the actions for HubSpot, and then choose **Next** again.
7. Specify who can access the connector:
- To share with a team, enter a team name or group email.
 - To keep it private, leave the field blank.
 - To share with all users, toggle **Everyone in your organization**.
8. Choose **Publish**.

The connector appears in the **Available** panel with a status of **Connected**.

Using available actions

After you set up the connector, you can use the following read-only actions to access your HubSpot CRM data.

HubSpot available actions

Action	Description
Get CRM Objects	Fetches multiple CRM objects of the same object type in a single request.

Action	Description
Get properties	Fetches properties for the CRM objects.
Get user details	Fetches user details.
Search CRM Objects	Searches across CRM objects.
Search Owners	Searches across owners.
Search properties	Searches across properties.

Note

This connector supports read-only actions. The actions that you can use depend on the scopes configured in your HubSpot app.

Troubleshooting

To edit, share, or delete your connector, see [the section called “Managing existing integrations”](#).

Resolving authentication issues

- **Sign-in fails** – Verify that your HubSpot account is active and that you can sign in to the [HubSpot sign-in page](#) on the HubSpot website. If your organization uses single sign-on (SSO), confirm that your identity provider is configured correctly.
- **Access denied** – Confirm that your HubSpot user has an admin role with permissions to access CRM data. Verify that the client ID and client secret from your app are correct and that the required read scopes are configured.

HuggingFace integration

With HuggingFace integration in Amazon Quick, you can access machine learning models, datasets, and spaces through MCP server connectivity. This integration provides action capabilities for ML workflow operations and model management.

What you can do

HuggingFace integration provides connector capabilities through MCP server connectivity:

- Browse and download models from HuggingFace Hub
- Access and manage datasets
- Interact with HuggingFace Spaces
- Upload and manage your own models
- Run inference on hosted models
- Manage model repositories and versions

Available tools

The HuggingFace MCP server typically provides these tools:

- `search_models` - Search for models on HuggingFace Hub
- `get_model_info` - Get detailed model information
- `download_model` - Download models locally
- `list_datasets` - List available datasets
- `get_dataset_info` - Get dataset information
- `run_inference` - Run inference on hosted models
- `upload_model` - Upload models to Hub
- `list_spaces` - List HuggingFace Spaces

Note

The specific tools and capabilities available through this MCP server may change over time. For the most current information about supported tools, features, and implementation details, check the official HuggingFace documentation and MCP server repository.

Setting up HuggingFace integration

HuggingFace integration uses MCP server connectivity to provide action capabilities. For detailed setup instructions, see [the section called “MCP integration”](#).

You'll need:

- HuggingFace account with appropriate permissions
- HuggingFace API token for authentication

Compatibility

HuggingFace integration supports:

- **Chat Agents:** Yes
- **Flows:** Yes
- **Knowledge Base:** No

Intercom integration

With Intercom integration in Amazon Quick, you can manage customer conversations, support tickets, and user engagement through MCP server connectivity. This integration provides action capabilities for customer support and communication operations.

What you can do

Intercom integration provides connector capabilities through MCP server connectivity:

- Manage customer conversations and messages
- Create and update support tickets
- Manage user profiles and contact information
- Send targeted messages and campaigns
- Track customer engagement and activity
- Manage team assignments and workflows

Available tools

The Intercom MCP server typically provides these tools:

- `create_conversation` - Start new conversations
- `reply_to_conversation` - Reply to existing conversations

- `list_conversations` - List customer conversations
- `create_user` - Create new user profiles
- `update_user` - Update user information
- `send_message` - Send messages to users
- `create_ticket` - Create support tickets
- `search_users` - Search for users

Note

The specific tools and capabilities available through this MCP server may change over time. For the most current information about supported tools, features, and implementation details, check the official Intercom documentation and MCP server repository.

Setting up Intercom integration

Intercom integration uses MCP server connectivity to provide action capabilities. For detailed setup instructions, see [the section called “MCP integration”](#).

You'll need:

- Intercom account with appropriate permissions
- Intercom API access token

Compatibility

Intercom integration supports:

- **Chat Agents:** Yes
- **Flows:** Yes
- **Knowledge Base:** No

Linear integration

With Linear integration in Amazon Quick, you can manage issues, track projects, and streamline development workflows through MCP server connectivity. This integration provides action capabilities for project management and issue tracking operations.

What you can do

Linear integration provides connector capabilities through MCP server connectivity:

- Create and manage issues and tasks
- Track project progress and milestones
- Manage team workflows and assignments
- Update issue status and priorities
- Create and manage project cycles
- Generate reports and track metrics

Available tools

The Linear MCP server typically provides these tools:

- `create_issue` - Create new issues
- `update_issue` - Update issue details
- `list_issues` - List team issues
- `search_issues` - Search for specific issues
- `create_project` - Create new projects
- `list_projects` - List team projects
- `assign_issue` - Assign issues to team members
- `create_cycle` - Create project cycles

Note

The specific tools and capabilities available through this MCP server may change over time. For the most current information about supported tools, features, and implementation details, check the official Linear documentation and MCP server repository.

Setting up Linear integration

Linear integration uses MCP server connectivity to provide action capabilities. For detailed setup instructions, see [the section called “MCP integration”](#).

You'll need:

- Linear account with appropriate team permissions
- Linear API key for authentication

Compatibility

Linear integration supports:

- **Chat Agents:** Yes
- **Flows:** Yes
- **Knowledge Base:** No

Local connectors

In addition to cloud connectors, you can add a local connector that runs on your own machine, such as a local MCP server. A local connector lets Amazon Quick use tools you host yourself, such as a database, an internal API, or a developer tool, so that those tools become available to agents in chat and in scheduled tasks. Because a local connector runs on your machine, its tools are available only when that machine is on and connected. For an overview of cloud connectors, see [Connectors](#).

To add a local connector, open the Desktop App, go to **Customize**, and choose the **Connectors** tab. Choose **Create**, and then choose **Create local connector**. A dialog appears with three connection types: **Local**, **Import**, and **Remote**. Choose the type that matches how your MCP server runs.

Local

Local runs an MCP server as a command on your machine. This is the most common option for development tools and locally installed servers.

Local connection fields

Field	Required	Description
Name	Yes	Enter a descriptive name to identify this MCP server, such as <code>My Database MCP</code> .
Command	Yes	Enter the executable to run, such as <code>python</code> , <code>npx</code> , <code>node</code> , or <code>uvx</code> .
Arguments	No	Enter the command line arguments, separated by spaces, such as <code>-m mcp_server --port 8080</code> .
Description	No	Describe what this MCP server does and what tools it provides. This description helps Amazon Quick understand when to use the server's tools.
Environment variables	No	Add the environment variables that the server requires. Choose + Add variable to add key-value pairs.
Timeout (seconds)	No	Enter the maximum time to wait for the server to start, from 5 to 300 seconds. The default is 30 seconds.

Import

Import loads MCP server configurations from an existing configuration file. This is useful when you already have MCP servers configured for another tool and want to reuse that configuration. When you enter a configuration file path, Amazon Quick scans the file, detects the MCP servers defined in it, and detects the tools installed on your system.

Import connection fields

Field	Required	Description
Config file path	Yes	Enter the path to the configuration file on your machine, such as <code>~/.kiro/settings/mcp.json</code> .

Remote

Remote connects to an MCP server running on a remote machine or cloud service over HTTP. This is useful for shared team servers or hosted MCP services.

Remote connection fields

Field	Required	Description
Name	Yes	Enter a descriptive name to identify this MCP server.
URL	Yes	Enter the URL of the MCP endpoint, such as <code>https://mcp.example.com/sse</code> .
Token	No	Enter a bearer token for authentication with the remote server.
Description	No	Describe what this MCP server does and what tools it provides.
Timeout (seconds)	No	Enter the maximum time to wait for the server to respond, from 5 to 300 seconds. The default is 30 seconds.

Microsoft OneDrive integration

Use the Microsoft OneDrive integration to perform actions on OneDrive files and folders, or to create knowledge bases from OneDrive content for AI-powered search and Q&A.

What you can do

Connector

Create, update, delete, and manage OneDrive files and folders through Microsoft Graph API calls.

Knowledge base

Index OneDrive documents, spreadsheets, and presentations. Amazon Quick agents can then search and answer questions about the content. Amazon Quick supports two setup methods:

user-managed setup for quick configuration, and admin-managed setup with service credentials for document-level access control.

Microsoft OneDrive action integration

Use the Microsoft OneDrive connector to manage OneDrive files, folders, and Excel workbooks directly in Amazon Quick through natural language.

Amazon Quick supports multiple authentication methods for Microsoft OneDrive. Choose the method that best fits your organization's security requirements.

- **Default OAuth app** – Uses an AWS-managed OAuth application. No additional credentials are needed. Users authenticate directly with their Microsoft account.
- **Custom OAuth app** – Uses a customer-managed application registered in Microsoft Entra. This option gives your organization full control over the OAuth configuration. Users authenticate on behalf of a signed-in user (delegated permissions).
- **Service-to-Service OAuth** – Uses client credentials for server-to-server authentication without user interaction (application permissions). Suitable for automated workflows.

For more information about the authentication methods that Amazon Quick supports, see [the section called “Authentication methods”](#).

Before you begin

Make sure you have the following before you set up the integration.

- A Microsoft 365 account with OneDrive access.
- For **Custom OAuth app** or **Service-to-Service OAuth**: Access to the [Microsoft Entra admin center](#) on the Microsoft website with at least Application Developer permissions.
- For subscription requirements, see [the section called “Set up integrations in the console”](#).

Configure Microsoft Entra

If you are using **Default OAuth app** authentication, skip this section and proceed to [the section called “Setting up the connector in Amazon Quick”](#).

Before you configure Amazon Quick, create an app registration in Microsoft Entra. Complete all of the following steps in Entra before moving to the Amazon Quick console.

For more information about app registrations, see [Register an application with the Microsoft identity platform](#) in the Microsoft documentation.

Register the application

1. Open the [Microsoft Entra admin center](#).
2. In the left navigation, choose **Entra ID**, then choose **App registrations**.
3. Choose **New registration**.
4. For **Name**, enter a descriptive name for your integration.
5. For **Supported account types**, choose **Accounts in this organizational directory only**.
6. For **Redirect URI**, select **Web** and enter `https://{region}.quicksight.aws.amazon.com/sn/oauthcallback`. Replace `{region}` with the AWS Region where your Amazon Quick instance is deployed.
7. Choose **Register**.
8. On the overview page, copy the **Application (client) ID** and **Directory (tenant) ID**. You need these values for the Amazon Quick configuration.

Create a client secret

Amazon Quick needs a client secret to authenticate with Microsoft Entra. This secret acts as a password for the app registration.

1. From your app registration, choose **Certificates & secrets**.
2. Choose **New client secret**.
3. Enter a description and choose an expiration period.
4. Choose **Add**.
5. Copy the **Value** immediately. This value is only displayed once.

Important

Copy the secret **Value**, not the Secret ID. The Value is the longer string used for authentication.

Configure API permissions

This integration uses delegated permissions, which allow the app to act on behalf of a signed-in user. For more information, see [Overview of Microsoft Graph permissions](#) in the Microsoft documentation.

1. From your app registration, choose **API permissions**.
2. Choose **Add a permission**, then choose **Microsoft Graph**.
3. Choose **Delegated permissions** and add the permissions from the table below.
4. Choose **Grant admin consent for [your tenant name]** to approve the permissions.

Add the following as Delegated permissions in your Entra app registration. For the full permissions reference, see [Microsoft Graph permissions reference](#) in the Microsoft documentation.

OneDrive action integration – delegated permissions

Permission	Description
Files.ReadWrite	Allows the app to read, create, update, and delete the signed-in user's files.
Files.ReadWrite.All	Allows the app to read, create, update, and delete all files the signed-in user can access.
offline_access	Allows the app to refresh access tokens without requiring the user to sign in again. This reduces how often users need to re-authenticate.
Sites.Read.All	Allows the app to read documents and list items in all site collections on behalf of the signed-in user.
User.Read	Allows the app to read the signed-in user's profile.
User.Read.All	Allows the app to read the full set of profile properties of all users in the organization on behalf of the signed-in user.

Record your credentials

Before leaving the Microsoft Entra admin center, confirm you have the following values. You need them for the Amazon Quick configuration.

Required credentials from Microsoft Entra

Value	Where to find it
Application (client) ID	App registration overview page
Directory (tenant) ID	App registration overview page
Client secret value	Certificates & secrets page

Setting up the connector in Amazon Quick

Connect from the Available tab

If you want to use Default OAuth app authentication, you can connect directly from the **Available** tab without additional configuration.

1. In the Amazon Quick console, choose **Connectors**.
2. On the **Available** tab, find **OneDrive** and choose **Connect**.
3. Complete the Microsoft sign-in flow and grant the requested permissions.

To configure a connector with Custom OAuth app or Service-to-Service OAuth instead, use the **Create for your team** tab as described below.

Create from the Create for your team tab

After you complete any required Entra configuration, create the connector in Amazon Quick.

1. In the Amazon Quick console, choose **Connectors**.
2. Choose the **Create for your team** tab.
3. Find and choose **Microsoft OneDrive**.

 Note

If a Microsoft OneDrive connector already exists, a dialog appears with your existing connectors. To use an existing connector, choose it. To create a new one, choose **No, create new**.

4. On the **Integration type** page, select **Perform actions in Microsoft OneDrive** and choose **Next**.
5. Enter a **Name** for your connector. Optionally, choose **+ Add Description** to add a description.
6. For **Connection type**, choose **Public network**.
7. For **OAuth Configuration**, choose one of the following authentication methods and configure the required fields.

- a. For **Default OAuth app**:

No additional credentials are needed. Choose **Next** to continue.

- b. For **Custom OAuth app** (user authentication with delegated permissions), configure the following fields:

- **Base URL** (Optional) – The Microsoft Graph API base URL. Example: `https://graph.microsoft.com/v1.0`
- **Client ID** – The Application (client) ID from your Entra app registration.
- **Client secret** – The client secret value from your Entra app registration.
- **Token URL** – The token endpoint. Example: `https://login.microsoftonline.com/{tenant-id}/oauth2/v2.0/token`
- **Authorization URL** – The authorization endpoint. Example: `https://login.microsoftonline.com/{tenant-id}/oauth2/v2.0/authorize`
- **Redirect URL** – Pre-filled with the Amazon Quick callback URL.

- c. For **Service-to-Service OAuth** (service authentication with application permissions), configure the following fields:

- **Base URL** (Optional) – The Microsoft Graph API base URL. Example: `https://graph.microsoft.com/v1.0`
- **Client ID** – The Application (client) ID from your Entra app registration.
- **Client secret** – The client secret value from your Entra app registration.

- **Token URL** – The token endpoint. Example: `https://login.microsoftonline.com/{tenant-id}/oauth2/v2.0/token`

 Note

The scope for the client credentials token request (`https://graph.microsoft.com/.default`) is set automatically by Amazon Quick. You do not need to configure it manually.

8. Choose **Next**.
9. If you chose **Default OAuth app** or **Custom OAuth app**, a Microsoft authorization window opens. Review the requested permissions and choose **Accept**.

If you see an error instead of the consent dialog, your organization might restrict third-party app access. See [the section called "Admin consent for Microsoft 365"](#).

10. On the **Review** page, review the available actions for the connector. Choose **Next**.
11. On the **Publish** page, choose who can access the connector. You can enable access for everyone in your organization or search for specific teams or groups.
12. Choose **Publish**.

Available actions

After you set up the integration, the following actions are available.

Microsoft OneDrive available actions

Category	Action	Description
Drive and items	Get Drive	Retrieve the properties and relationships of a drive.
Drive and items	List Item	Get all items contained in the drive.
Drive and items	Get Item	Get an item contained in the drive.

Category	Action	Description
Drive and items	Create Folder	Create a new folder in the user's drive.
Drive and items	Update Item	Update the metadata for a file or folder.
Drive and items	Delete Item	Delete a file or folder. Moves the item to the recycle bin.
Drive and items	List Child Folders	Return a collection of items in the children of a folder.
Drive and items	Copy Item	Create a copy of a file or folder to another location.
Drive and items	Add Permissions	Send a sharing invitation for a file or folder.
Drive and items	Upload File	Upload a new file or update an existing file. Supports files up to 250 MB.
Excel workbooks	List Sheets	Retrieve a list of worksheet objects.
Excel workbooks	Add Sheet	Add a new worksheet to the workbook.
Excel workbooks	Read Sheet	Retrieve the properties of a worksheet object.
Excel workbooks	Update Sheet	Update the properties of a worksheet object.
Excel workbooks	Delete Sheet	Delete the worksheet from the workbook.
Excel workbooks	Read Cell	Get the value of a single cell by row and column number.

Category	Action	Description
Excel workbooks	Write Cell	Set the value of a single cell by row and column number.
Excel workbooks	Read Range	Get the values of a range.
Excel workbooks	Write Range	Update the values of a range.
Excel workbooks	Clear Range	Clear range values, format, fill, and border.
Excel workbooks	Delete Range	Delete the cells associated with the range.
Excel workbooks	Get Used Range	Get the smallest range that encompasses cells with a value or formatting.

Manage and troubleshoot

To edit, share, or delete your integration, see [the section called “Managing existing integrations”](#).

Authentication issues

- **Incorrect app registration** – Verify the app registration in Microsoft Entra includes the required API permissions and that admin consent has been granted.
- **Expired client secret** – Check if the client secret has expired in **Certificates & secrets** and generate a new one if needed.
- **Incorrect redirect URI** – Verify the redirect URI in Microsoft Entra matches `https://{region}.quicksight.aws.amazon.com/sn/oauthcallback`.

Common error messages

- **Access denied. You do not have permission to perform this action** – The authenticated user does not have the required permissions. Contact your administrator to verify and grant appropriate permissions.

- **AADSTS50020: User account from identity provider does not exist in tenant**
 - The user account is not configured in the correct Microsoft Entra tenant. Verify the user account exists in the tenant that matches the Directory (tenant) ID in your app registration.

Admin consent for Microsoft 365

When you use the **Default OAuth app** authentication method, Amazon Quick uses an AWS-managed application to access Microsoft OneDrive on behalf of the signed-in user. Most users can complete setup without any extra steps. However, if your Microsoft 365 tenant restricts third-party app access, a Microsoft 365 administrator must grant one-time consent before users can connect.

If you see an error when you sign in during connector setup, your organization might restrict third-party app access. Share the following information with your Microsoft 365 administrator:

- **What to do:** Grant admin consent for the Amazon Quick Microsoft OneDrive integration application.
- **Why:** Amazon Quick needs delegated access to OneDrive files, folders, and Excel workbooks to perform actions on behalf of users.

An administrator can grant consent in one of the following ways:

- **Through the consent dialog** – A Global Administrator or Privileged Role Administrator initiates the connector setup flow. In the Microsoft sign-in dialog, they select the **Consent on behalf of your organization** check box and choose **Accept**.
- **Through the Microsoft Entra admin center** – Sign in to the [Microsoft Entra admin center](#) on the Microsoft website. Choose **Enterprise applications**, locate the Amazon Quick application, choose **Permissions**, and choose **Grant admin consent for *Your Organization***.

After consent is granted, any user in your organization can connect without being prompted for individual consent.

Note

To check whether your tenant restricts user consent, go to the Microsoft Entra admin center and choose **Enterprise applications, Consent and permissions, User consent settings**. If

the setting is **Do not allow user consent**, an administrator must grant consent before users can use the connector.

Microsoft OneDrive knowledge base integration

Microsoft OneDrive is a cloud storage service for storing, sharing, and collaborating on files across your organization. With the Amazon Quick OneDrive connector, you can create knowledge bases from your OneDrive content to power AI-driven search and Q&A.

Amazon Quick supports two authentication methods for connecting to OneDrive:

- **User-managed setup** – You sign in to OneDrive directly to authorize the connection. This is the simplest way to get started. For more information, see [the section called “User-managed setup”](#).
- **Admin-managed setup (service credentials)** – An Entra ID app registration authenticates on behalf of the application using certificate-based credentials, and Amazon Quick crawls the OneDrive content of every user in your organization. Admin-managed setup includes built-in document-level access control (ACL). Amazon Quick automatically syncs access control lists from OneDrive and verifies each user's permissions at query time. Users see answers only from documents that they are authorized to access. For more information, see [the section called “Admin-managed setup”](#).

After connecting, Amazon Quick indexes your OneDrive files and folders into a knowledge base. Your Amazon Quick agents can then search this content and generate answers grounded in your OneDrive data.

Prerequisites

Before you set up the OneDrive knowledge base integration, make sure that you have the following:

- An AWS account with an active Amazon Quick instance.
- A Microsoft 365 account with OneDrive access.
- For subscription requirements, see [the section called “Set up integrations in the console”](#).

For admin-managed setup, additional prerequisites apply. For more information, see [the section called “Admin-managed setup”](#).

Supported content types

The OneDrive connector supports the common file types that Amazon Quick knowledge bases support, including the following:

- **Microsoft Office documents:** Word, Excel, PowerPoint
- **PDF files**
- **Text files and rich text documents**
- **Text documents with embedded images**
- **Audio and video files**

For more information about supported file types, size limits, and content processing options, see [Common configuration settings](#).

User-managed setup

With user-managed setup, you sign in to OneDrive directly and Amazon Quick handles authentication through a managed OAuth flow. Amazon Quick uses a pre-registered multi-tenant application, so you don't need to create an app registration. Most users can complete setup in a few minutes.

Prerequisites

Before you begin, make sure you have the following:

- A Microsoft 365 account with access to the OneDrive content you want to index.
- Access to Amazon Quick with permissions to create knowledge bases.
- A browser that allows popups from the Amazon Quick console domain.

Your Microsoft administrator might need to grant organizational consent before users can create a OneDrive knowledge base. Administrators can grant organization-wide consent by signing in and choosing **Consent on behalf of your organization** during the integration creation flow. For more information, see [the section called "Admin consent for Microsoft 365"](#).

Important

User-managed setup does not support document-level access control (ACL). All indexed content is accessible to any user who has access to the knowledge base in Amazon Quick.

Individual permissions in OneDrive are not enforced. Carefully review which content you include when creating a knowledge base. If you require document-level access control, use [Admin-managed setup \(service credentials\)](#) instead.

Create the OneDrive knowledge base

Navigate to Knowledge

1. In Amazon Quick, choose **Knowledge** from the left navigation panel.
2. Under **Set up new knowledge base**, locate **Microsoft OneDrive** and choose the **+** icon.

Sign in to OneDrive

The **Create OneDrive knowledge base** wizard opens on the Authentication method step. **Quick setup** is selected by default.

1. Under **Quick setup**, choose **Sign in to OneDrive**.
2. A Microsoft sign-in window opens. Enter your Microsoft 365 credentials.
3. If a permissions consent dialog appears, review the permissions and choose **Accept**.

If you see an error instead of the consent dialog, your organization might restrict third-party app access. See [the section called "Admin consent for Microsoft 365"](#).

4. After successful authentication, choose **Next**.

Choose content

1. Enter a **Name** and optional **Description** for your knowledge base.
2. In the **Content** section, choose **Add content**.
3. A dialog opens showing OneDrive files and folders accessible to your account. Select the files or folders you want to index.
4. Choose **Add** to confirm your selections. You can add more content or remove items before continuing.
5. Choose **Create**.

Initial sync

After you choose **Create**, you are returned to the knowledge base list page. The knowledge base might take a few minutes to finish provisioning. Once creation is complete, an initial sync is automatically triggered.

Admin consent for Microsoft 365

Most users complete setup without any extra steps. However, if your Microsoft 365 tenant restricts third-party app access, you might see an error when you sign in. In this case, a Microsoft 365 administrator needs to grant one-time consent for the Amazon Quick application. After consent is granted, any user in your organization can connect.

If you are not a Microsoft 365 administrator, share the following information with your administrator:

- **What to do:** Grant admin consent for the Amazon Quick OneDrive integration application.
- **Why:** Amazon Quick needs delegated read access to OneDrive files to index content for knowledge bases.

For detailed instructions on granting admin consent through the consent dialog or the Microsoft Entra admin center, see [the section called "Grant organization-wide admin consent"](#).

Permissions requested

The following delegated permissions are requested when a user signs in. Share this list with your administrator if they need to review the permissions before granting consent.

User-managed setup – permissions

Permission	API	Type	Description
Files.Read.All	Microsoft Graph	Delegated	Read all files the signed-in user can access.
Notes.Read.All	Microsoft Graph	Delegated	Read all OneNote notebooks the signed-in user can access.
User.Read	Microsoft Graph	Delegated	Sign in and read the user's profile.

Permission	API	Type	Description
offline_access	Microsoft Graph	Delegated	Maintain access using refresh tokens.

Manage and troubleshoot user-managed connections

To edit, share, or delete your integration, see [the section called “Managing existing integrations”](#).

- **App blocked by administrator** – Your Microsoft 365 tenant restricts third-party app access. Ask your Microsoft 365 administrator to grant admin consent. For more information, see [the section called “Admin consent for Microsoft 365”](#).
- **Sign-in window closes without completing** – Verify that your browser allows popups from the Amazon Quick console domain and that third-party cookies are enabled.
- **Token expired** – Delegated credentials last approximately 90 days. If syncs fail after this period, you need to re-authenticate. For more information, see [the section called “Token refresh with user-managed setup”](#).
- **Missing content** – Verify that the account you used for authentication has access to the files and folders you selected. Content shared with you after the initial sync requires a resync to be indexed.

For additional troubleshooting, including sync monitoring and reports, see [the section called “Troubleshooting”](#).

Known limitations

- Document-level access control (ACL) is not supported with user-managed setup. If you require document-level access control, use [Admin-managed setup \(service credentials\)](#).
- File comments synchronization is not supported.

Next steps

After your knowledge base is created and the initial sync completes, you can use it in Amazon Quick to answer questions from your OneDrive content. To manage your knowledge base, see [the section called “Managing existing integrations”](#).

Admin-managed setup (service credentials)

With admin-managed setup, an administrator configures an AWS KMS signing key and a Microsoft Entra ID app registration with certificate-based credentials. Individual users don't need to authorize through sign-in. Amazon Quick crawls the OneDrive content of every user in your organization.

Admin-managed setup includes built-in document-level access control list (ACL) support. Because admin-managed setup crawls the OneDrive content of all users, ACL is always enabled and cannot be turned off. Amazon Quick syncs ACLs from OneDrive during each crawl and verifies each user's permissions at query time, so users see answers only from documents that they are authorized to access. For more information, see [the section called "Document-level access controls"](#).

Prerequisites

Before you begin, make sure that you have the following:

- Administrator access to the Amazon Quick admin console.
- Administrator access to Microsoft Entra ID to register an application and grant API permissions.
- A Microsoft 365 tenant with OneDrive content to index.

Setup overview

The setup involves the following phases:

1. **Set up service credentials** – Create a KMS signing key, generate a certificate, register an application in Entra, and grant Amazon Quick permission to use the key. For more information, see [the section called "Set up service credentials"](#).
2. **Create the knowledge base in Amazon Quick** – Create a OneDrive knowledge base using the service credentials from Phase 1. For more information, see [the section called "Create the knowledge base"](#).

Document-level access control is automatically enabled for all admin-managed knowledge bases. For more information about how access controls work, see [the section called "Document-level access controls"](#).

Set up service credentials

Before you create the knowledge base in Amazon Quick, complete the following configuration steps in AWS and Microsoft Entra ID. You create a KMS signing key, generate a certificate, register an application in Entra, and grant Amazon Quick permission to use the key.

This setup involves multiple systems and might require coordination between different administrators in your organization. The following table summarizes each step and the role needed to complete it.

Setup steps and roles

Step	What you do	Role needed
1. KMS key	Create an asymmetric signing key in AWS KMS.	AWS administrator (KMS and IAM console access)
2. Certificate	Generate a self-signed certificate using the KMS public key.	Same as Step 1 (AWS CLI and OpenSSL required)
3. Entra app	Register an application in Microsoft Entra, assign API permissions, and upload the certificate.	Microsoft 365 Global Admin or Privileged Role Admin
4. KMS key access	Grant Amazon Quick permission to use the KMS key for signing.	Amazon Quick administrator (Admin Pro)
5. Create KB	Create the knowledge base in Amazon Quick using the credentials from the previous steps.	Any Amazon Quick user (Author Pro or Admin Pro)

Tip

In many organizations, a single person with both AWS and Microsoft 365 administrator access can complete all steps. If responsibilities are split across teams, share this table to coordinate the setup.

Prerequisites

Before you begin, make sure that you have the following:

- An AWS account with an active Amazon Quick instance.
- Access to the AWS KMS console (for creating the signing key).
- Amazon Quick administrator access (Admin Pro role) for granting KMS key permissions.
- A Microsoft 365 tenant with OneDrive.
- Global Administrator or Privileged Role Administrator access in Microsoft Entra ID.
- OpenSSL 3.0 or later and AWS CLI installed locally.
- The AWS account and Amazon Quick instance must be in the same Region.

Permissions

Admin-managed setup crawls the OneDrive content of every user in your organization and always enforces document-level access control. The Entra app registration requires the following Microsoft Graph application permissions.

Admin-managed setup permissions

Permission	Type	Purpose
Files.Read.All	Application	Read files in all users' OneDrive content.
Sites.Read.All	Application	Enumerate and read the OneDrive drives across the tenant. OneDrive for Business is hosted on SharePoint, so this permission is required to access each user's drive.
User.Read.All	Application	Read user profiles to resolve document-level access control.
Group.Read.All	Application	Read group objects to resolve group-based access control.
GroupMember.Read.All	Application	Read group memberships to resolve document-level access control.

 Important

Select the **Application permissions** tab, not **Delegated permissions**, when you assign these permissions in Entra. Admin-managed setup uses the client credentials flow, which requires application permissions.

 Note

OneNote crawling (Notes.Read.All) is not supported in admin-managed setup. Microsoft retired app-only tokens for OneNote APIs on March 31, 2025. Use [User-managed setup](#) for OneNote content.

Values collected during setup

The following table summarizes the values you create or collect during setup and where you use them.

Values reference

Value	Created in step	Used in step
KMS Key ARN	1 (KMS)	2 (Certificate), 4 (IAM), Quick setup
Certificate file (certificate.cer)	2 (Certificate)	3 (Entra upload)
Certificate thumbprint (base64url)	2 (Certificate)	Quick setup
Application (Client) ID	3 (Entra)	Quick setup
Directory (Tenant) ID	3 (Entra)	Quick setup

Step 1: Create an AWS KMS asymmetric signing key

Amazon Quick uses an AWS KMS asymmetric key to sign OAuth assertions when authenticating with Microsoft Entra ID. The private key never leaves KMS. Only the public key is exported and embedded in a certificate that gets uploaded to your Entra app registration.

Create the KMS key

1. Open the [AWS KMS console](#).
2. In the left navigation, choose **Customer managed keys**.
3. Choose **Create key**.

Configure the key

On the **Configure key** page, set the following values:

KMS key configuration

Setting	Value
Key type	Asymmetric
Key usage	Sign and verify
Key spec	RSA_2048
Key material origin	KMS (recommended)
Regionality	Single-Region key (default). Multi-Region keys are not supported.

Add labels

On the **Add labels** page, enter an alias for the key. For example: quick-onedrive-service-auth.

Note

The key administrator and key usage permissions on the following pages are optional. The defaults are sufficient for this setup. You grant Amazon Quick access to the key separately in Step 4.

Choose **Skip to review**, then choose **Finish** to create the key.

Record the Key ARN

After the key is created, open the key detail page and record the **Key ARN**. The ARN has the following format:

```
arn:aws:kms:us-west-2:123456789012:key/xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxxxx
```

You need this value in Steps 2, 4, and when you create the knowledge base in Quick.

Step 2: Generate a self-signed certificate

Microsoft Entra ID requires an X.509 certificate to validate signed assertions. Because the KMS private key never leaves AWS KMS, you cannot use it directly with OpenSSL. Instead, you generate a temporary local key pair and create a certificate signing request. Then, use the OpenSSL -force_pubkey option to inject the KMS public key into the final certificate. The result is a self-signed certificate whose public key matches the KMS key pair.

Prerequisites

- AWS CLI installed and configured.
- OpenSSL 3.0 or later.
- The KMS Key ARN from Step 1.

Generate the certificate

Run the following commands in a terminal. Replace the *placeholder* values with your own.

Verify OpenSSL version

```
openssl version
```

Confirm the output shows version 3.0 or later.

Export the KMS public key

```
aws kms get-public-key \  
  --key-id KMS_KEY_ARN \  
  --region REGION \  
  --output text \  
  --query PublicKey | base64 --decode > public_key.der
```

Note

On macOS, use `base64 --decode` or `base64 -D` depending on your shell environment.

Convert the public key to PEM format

```
openssl rsa -pubin -inform DER -in public_key.der -outform PEM -out kms_public_key.pem
```

Generate a temporary local key pair

```
openssl genrsa -out temp_private_key.pem 2048
```

Create a certificate signing request

```
openssl req -new \  
  -key temp_private_key.pem \  
  -out cert.csr \  
  -subj "/CN=QuickOneDriveServiceAuth/O=YourOrganization/C=US"
```

Generate the certificate with the KMS public key

```
openssl x509 -req \  
  -in cert.csr \  
  -signkey temp_private_key.pem \  
  -out certificate.pem \  
  -days 730 \  
  -force_pubkey kms_public_key.pem
```

Note

OpenSSL displays the warning `Signature key and public key of cert do not match`. This is expected because the certificate is signed with the temporary local key but contains the KMS public key. The certificate is valid and works correctly with Microsoft Entra.

Convert to DER format for Entra upload

```
openssl x509 -in certificate.pem -outform DER -out certificate.cer
```

Clean up temporary files

```
rm -f temp_private_key.pem cert.csr public_key.der kms_public_key.pem certificate.pem
```

Important

Keep the `certificate.cer` file. You upload it to Microsoft Entra ID in Step 3.

Calculate the certificate thumbprint

Run the following command to calculate the base64url-encoded SHA-1 thumbprint of the certificate:

```
openssl dgst -sha1 -binary certificate.cer | base64 | tr '+/' '-_' | tr -d '='
```

Record this value. You enter it when you create the knowledge base in Quick.

Note

The base64url-encoded thumbprint is different from the hexadecimal thumbprint shown in the Microsoft Entra portal. Quick requires the base64url format.

Step 3: Register an application in Microsoft Entra ID

Register the application

1. Sign in to the [Microsoft Entra admin center](#) on the Microsoft website.
2. In the left navigation, expand **Entra ID** and choose **App registrations**.
3. Choose **New registration**.
4. For **Name**, enter QuickOneDriveServiceAuth.
5. For **Supported account types**, select **Accounts in this organizational directory only (Single tenant)**.
6. Leave **Redirect URI** blank. A redirect URI is not required because the application uses the client credentials flow, not an interactive sign-in flow.
7. Choose **Register**.

Record application details

On the application **Overview** page, record the following values:

Application details

Value	Location
Application (client) ID	Shown on the Overview page under Essentials .
Directory (tenant) ID	Shown on the Overview page under Essentials .

Configure API permissions

Add the Microsoft Graph application permissions listed in the [the section called "Permissions"](#) section.

1. In the left navigation of your app registration, choose **API permissions**.
2. Choose **Add a permission**.
3. Choose **Microsoft Graph**.
4. Choose **Application permissions**.
5. Search for and select `Files.Read.All`, `Sites.Read.All`, `User.Read.All`, `Group.Read.All`, and `GroupMember.Read.All`, then choose **Add permissions**.

⚠ Important

Select the **Application permissions** tab, not **Delegated permissions**. Admin-managed setup uses the client credentials flow, which requires application permissions.

Grant admin consent

1. On the **API permissions** page, choose **Grant admin consent for [Your Organization]**.
2. Confirm the consent when prompted.

⚠ Important

Admin consent is required for application permissions. Without it, the application cannot access OneDrive data.

Upload the certificate

1. In the left navigation of your app registration, choose **Certificates & secrets**.
2. Choose the **Certificates** tab.
3. Choose **Upload certificate**.
4. Select the `certificate.cer` file you generated in Step 2.
5. Choose **Add**.

ℹ Note

The Entra portal displays the certificate thumbprint in hexadecimal format. This is different from the base64url-encoded thumbprint you calculated in Step 2. Use the base64url value when you configure the knowledge base in Quick.

Step 4: Grant Amazon Quick permission to the KMS key

Amazon Quick needs permission to use the KMS key for signing OAuth assertions. You grant this permission from the Amazon Quick admin console.

Note

This step requires Amazon Quick administrator access (Admin Pro role). If you are not an administrator, ask your Amazon Quick admin to complete this step using the KMS key ARN from Step 1.

Important

If your organization manages its own Amazon Quick IAM service role, the following console steps might not apply. Instead, ensure the role has `kms:Sign` permission on the KMS key ARN from Step 1.

1. In Amazon Quick, choose **Manage account** from the left navigation pane.
2. Under **Permissions**, choose **AWS resources**.
3. On the AWS resources page, scroll to **AWS Key Management Service** and select the checkbox.
4. Choose **Select keys**.
5. In the **Select KMS keys** dialog, enter the KMS key ARN you recorded in Step 1 and choose **Add**.
6. The key ARN appears in the list. Choose **Finish**.
7. Choose **Save** at the bottom of the AWS resources page.

Next steps

After you complete the setup, create the OneDrive knowledge base connection in Amazon Quick. For instructions, see [the section called "Create the knowledge base"](#).

Create the knowledge base in Amazon Quick

After you complete the service credentials setup, create the knowledge base in Amazon Quick and provide the credentials. Any enterprise user can complete this step. Amazon Quick administrator access is not required.

You need the following values from [the section called "Set up service credentials"](#). If an administrator completed the setup on your behalf, collect these values before you proceed:

- KMS key ARN (from Step 1)

- Certificate thumbprint in base64url format (from Step 2)
- Entra application (client) ID (from Step 3)
- Directory (tenant) ID (from Step 3)

Navigate to Knowledge

1. In Amazon Quick, choose **Knowledge** from the left navigation pane.
2. Under **Set up new knowledge base**, locate **Microsoft OneDrive** and choose the **Add** icon.

Select service credentials authentication

The **Create OneDrive knowledge base** wizard opens on the **Authentication method** step. Choose **Have admin access? Connect with service credentials to enable document-level access control**.

If you have previously created a knowledge base using service credentials, the **Connected account** dropdown appears with your most recent connection pre-selected. When you select an existing connection, the Tenant ID and Client ID are displayed below the dropdown for verification. You can use an existing connection or choose **+ Add account** to enter new service credentials.

Enter service credentials

After choosing **+ Add account**, fill in the following fields:

Service credentials form

Field	Description	Example
Name	A descriptive name for this OneDrive connection. You see this name when selecting a connection for future knowledge bases.	OneDrive Production
Tenant ID	The Microsoft Entra (Azure AD) tenant identifier.	34567890-1234-1234-1234-123456789012
Client ID	The Application (client) ID from your Entra app registration.	12345678-1234-4321-4321-210987654321

Field	Description	Example
KMS key ARN	The ARN of the AWS KMS asymmetric signing key.	arn:aws:kms:us-west-2:123456789012:key/my-key-id
Certificate Thumbprint	The base64url-encoded SHA-1 thumbprint of the certificate uploaded to Entra.	abc123DEF456ghi789 EXAMPLE

Fill in each field and choose **Next**.

 Tip

Give the connection a recognizable name. This name appears in the **Connected account** dropdown when you create future knowledge bases.

Choose content

Enter a **Name** and optional **Description** for your knowledge base.

Admin-managed setup crawls the OneDrive content of every user in your organization. Under **All users**, Amazon Quick syncs files and permissions for all users. You do not select individual sites, drives, or paths.

Choose **Next**.

 Note

Because admin-managed setup syncs all users' OneDrive content, the Entra app registration must have tenant-wide Sites.Read.All permission. For the full permission set, see [the section called "Permissions"](#).

 Note

Folders shared with a user are not crawled in admin-managed setup. Shared folders are stored in SharePoint and cannot be accessed with application credentials. To sync shared

content, create a SharePoint knowledge base instead. For more information, see [the section called "Known limitations"](#).

Configure additional settings

In the final step of the wizard, you can configure optional settings for the knowledge base.

Note

Document-level access control (ACL) is always enabled for admin-managed OneDrive knowledge bases and cannot be turned off. When you enable ACL by creating an admin-managed knowledge base, the console might display a note that one-time admin consent is required for real-time access verification. If you are a Microsoft 365 administrator, choose the **Grant admin consent** link to grant consent directly. If you are not an administrator, ask your Microsoft 365 admin to grant consent. For more information, see [the section called "Admin consent"](#).

Filter content by date

Limits which documents are crawled based on their last modified date in OneDrive. The start date defaults to one year before today. You can change or clear the start date, and optionally set an end date.

Multi-media content, file size, and file patterns

- **Visual content in documents** – Extracts and indexes visual elements from supported document formats. Enabled by default.
- **Audio files** – Transcribes and indexes audio files.
- **Video files** – Transcribes and indexes video files.

Note

Enabling audio and video indexing increases sync time and storage usage. Enable these options only if your OneDrive content includes relevant media files.

Choose **Create** to create the knowledge base.

Initial sync

After you choose **Create**, you are returned to the knowledge base list page. The knowledge base might take a few minutes to finish provisioning. Once creation is complete, an initial sync is automatically triggered. You do not need to start it manually.

Document-level access controls

Admin-managed OneDrive knowledge bases include built-in document-level access control. Because admin-managed setup crawls the OneDrive content of every user in your organization, ACL is always enabled and cannot be turned off. Amazon Quick syncs access control lists (ACLs) from OneDrive during each crawl. The system verifies each user's permissions at query time, so users only see answers from documents they are authorized to access in OneDrive.

How it works

When a user queries an Amazon Quick agent that uses an admin-managed OneDrive knowledge base, the system enforces access controls in two stages:

1. **Pre-retrieval filtering** – Amazon Quick performs a semantic search against the vector index to find the most relevant document passages. The system applies access control lists that were synced from OneDrive during the last crawl. This produces a preliminary set of candidate documents.
2. **Real-time verification** – The system verifies the candidate documents in real time by checking the querying user's current access in OneDrive. Only documents the user is currently authorized to access are included in the response.

This two-stage approach provides document-level access control that stays current even when OneDrive permissions change between syncs.

ACL management

Document-level access control is automatically enabled for all admin-managed OneDrive knowledge bases. No additional configuration is required, and the setting cannot be turned off.

For ACL enforcement to work, your Entra app registration must have the following Microsoft Graph application permissions:

- `User.Read.All`, `Group.Read.All`, and `GroupMember.Read.All` to resolve user and group membership.

- `Sites.Read.All` to enumerate and read the OneDrive drives, which are hosted on SharePoint.

For more information about ACL best practices, see [the section called “Best practices for managing ACLs in knowledge bases”](#).

Real-time access verification

The real-time verification stage uses a delegated OAuth flow managed automatically by Amazon Quick. Quick creates and manages a separate Microsoft Entra application specifically for this purpose. No customer configuration is required for this app. It is distinct from both the admin-managed app registration you created during setup and any user-managed OAuth app.

1. A user asks a question in the Quick chat assistant.
2. If the answer involves OneDrive content from an admin-managed knowledge base, Quick prompts the user to **Sign in to OneDrive**.
3. The user signs in and accepts the Microsoft consent dialog (if admin consent has not been granted).
4. Quick uses the user's delegated token to verify access to each candidate document in real time.
5. Only documents the user currently has access to in OneDrive are included in the response.

The sign-in is a one-time step. The delegated credentials use a refresh token and last approximately 90 days.

Delegated permissions

The real-time ACL app requests the following delegated permissions:

Real-time ACL – delegated permissions

Permission	Scope	Purpose
Read items in all site collections	<code>Sites.Read.All</code>	Verify user access to OneDrive content hosted on SharePoint.
Have full access to your files	<code>Files.ReadWrite.All</code>	Verify user access to specific files.
View your basic profile	<code>User.Read</code>	Identify the signed-in user.

Permission	Scope	Purpose
Maintain access to data you have given it access to	offline_access	Refresh tokens so users don't need to re-authenticate frequently.

Admin consent

The real-time ACL check uses a separate Microsoft Entra application from the one used in user-managed setup or the admin-managed app registration. If your organization requires admin consent, an administrator must grant consent for each application independently.

When you create an admin-managed knowledge base, the Amazon Quick console provides a direct link to grant admin consent. This link is for the real-time ACL application. If you are a Microsoft 365 administrator, you can grant consent directly from the console. Otherwise, share the link with your administrator.

If admin consent is not granted, each user sees the consent dialog on their first query that involves OneDrive content. After accepting, they are not prompted again for approximately 90 days.

For detailed instructions on granting admin consent through the consent dialog or the Microsoft Entra admin center, see [the section called “Grant organization-wide admin consent”](#).

Next steps

For more information about ACL best practices, see [the section called “Best practices for managing ACLs in knowledge bases”](#). For information about creating admin-managed OneDrive knowledge bases, see [the section called “Admin-managed setup”](#).

Known limitations

- OneNote files are not crawled in admin-managed setup. Microsoft retired app-only tokens for the OneNote APIs on March 31, 2025. To index OneNote content, use [User-managed setup](#).
- Shared folders are not crawled. Folders shared with a user are stored in SharePoint and cannot be accessed with the application credentials used for admin-managed setup. To sync shared content, we recommend creating a [Microsoft SharePoint knowledge base integration](#).

Manage and troubleshoot admin-managed connections

To edit, share, or delete your integration, see [the section called “Managing existing integrations”](#).

- **Unable to access KMS key** – Verify the KMS key ARN and Region. Confirm the KMS key has been added in the Amazon Quick admin console under **Manage account, AWS resources, AWS Key Management Service**. Confirm the key is enabled and has not been scheduled for deletion. Multi-Region keys are not currently supported.
- **Certificate validation failed** – Verify the thumbprint using the base64url-encoded SHA-1 value from the certificate generation step. Ensure the certificate uploaded to Entra has not expired.
- **ACL not enforced** – Confirm the Entra app has `Files.Read.All`, `Sites.Read.All`, `User.Read.All`, `Group.Read.All`, and `GroupMember.Read.All` on Microsoft Graph. Re-run a full sync after fixing permissions. For more information about verifying document access, see [the section called “Check document access \(ACL verification\)”](#).
- **Zero items crawled** – The sync completed but no documents were indexed. This typically indicates a permissions issue. Verify the Entra app has the correct API permissions, including `Sites.Read.All` so that Amazon Quick can enumerate the OneDrive drives hosted on SharePoint. Also verify that users have OneDrive content and that admin consent has been granted.
- **Syncs failing after certificate expiry** – If syncs fail across multiple knowledge bases that share the same data source connection, the certificate uploaded to Entra might have expired. Generate a new certificate (see [the section called “Step 2: Generate a self-signed certificate”](#)), upload it to the Entra app registration, and update the connection details. An Amazon Quick administrator can reassign data source ownership through **Manage assets** if the original creator is unavailable. For more information, see [the section called “Sharing data source connections”](#).
- **No results from ACL-enabled knowledge base** – If users receive no results from an admin-managed knowledge base, admin consent for the real-time ACL application might not have been granted. Your tenant might also block user consent. Grant admin consent using the link provided in the Amazon Quick console, or see [the section called “Admin consent”](#).

For additional troubleshooting, including sync monitoring, reports, and ACL verification, see [the section called “Troubleshooting”](#).

Troubleshooting OneDrive knowledge bases

Use this section to diagnose and resolve common issues with OneDrive knowledge base integrations.

Sign-in errors with user-managed setup

User-managed setup sign-in errors

Issue	Cause	Resolution
"Need admin approval" error during sign-in	Your tenant restricts user consent for third-party apps.	Ask a Global Administrator to grant admin consent. For more information, see the section called "Admin consent for Microsoft 365" .
Sign-in window closes without completing	Browser popup blocker or cookie settings.	Allow popups from the Amazon Quick domain and ensure third-party cookies are enabled.
"AADSTS65004: User declined to consent"	User chose Cancel or Deny on the consent dialog.	Retry the sign-in and choose Accept on the consent dialog.

Token refresh with user-managed setup

User-managed setup uses the `offline_access` scope to obtain refresh tokens. Delegated credentials last approximately 90 days. If the refresh token expires or is revoked (for example, by a password change, session revocation, or admin action), syncs fail and you need to re-authenticate.

To re-authenticate:

1. Navigate to the **Knowledge bases** list page in Amazon Quick.
2. Locate the knowledge base that requires a refresh. In the **Actions** column, choose the three-dot menu and choose **Edit integration**.
3. The **Edit integration** dialog opens. You can optionally update the integration name.
4. Choose **Save and Reauthenticate**.
5. A Microsoft sign-in window opens. Complete the sign-in flow using the same credentials that were originally used to create the integration.

⚠ Important

When you re-authenticate, sign in using the same Microsoft 365 account that was used during the original setup. Using different credentials might cause issues with connected resources. The new account might have different OneDrive permissions than the original account.

Common errors**Common OneDrive knowledge base errors**

Issue	Cause	Resolution
Unable to access KMS key	The KMS key ARN is incorrect, the key does not exist in the specified Region, or the KMS key has not been added in the Amazon Quick admin console.	Verify the KMS key ARN and Region. Confirm the KMS key has been added under Manage account, AWS resources , AWS Key Management Service . Confirm the key is enabled and has not been scheduled for deletion. Multi-Region keys are not supported.
Invalid Tenant ID	The Tenant ID is malformed or does not match a valid Entra tenant.	Copy the Tenant ID directly from the Entra admin center: Microsoft Entra ID, Overview .
Invalid Client ID	The Client ID is incorrect or the app registration does not exist.	Copy the Application (Client) ID from App Registrations , your app, Overview .
Certificate validation failed	The thumbprint does not match the certificate in Entra, or the certificate has expired.	Verify the thumbprint using the base64url-encoded SHA-1 value from the certificate generation step. Ensure the certificate has not expired.
Unable to connect to OneDrive	One or more credentials are incorrect, or the Entra app is missing	Review all field values. Confirm the Entra app has the correct API permissions.

Issue	Cause	Resolution
	required API permissions or admin consent.	ns and that admin consent has been granted.
Error loading the file picker with user-managed setup	The file picker is unable to establish a connection to OneDrive on behalf of the signed-in user.	Close the error message and choose Add content again to retry. If the error persists, close the creation dialog, reopen the flow, and sign in to OneDrive again.

Check document access (ACL verification)

Admin-managed OneDrive knowledge bases always enforce document-level access control. You can verify document-level access control for individual items using the Permission Checker in sync reports. For detailed instructions, see [the section called "Check document access \(ACL verification\)"](#).

Sync monitoring and reports

For detailed information about sync schedules, sync activity, sync reports, filtering items, and downloading reports, see [the section called "Sync reports and observability"](#).

Tip

To trigger a sync on demand, choose **Sync now** on the knowledge base detail page.

Next steps

After you resolve the issue, verify your fix by running a sync. For more information about sync monitoring, see [the section called "Sync reports and observability"](#).

Microsoft OneNote integration

With the Microsoft OneNote connector, you can manage notebooks, sections, and pages from your Microsoft 365 account directly in Amazon Quick through natural language. The connector uses Microsoft Graph for OneNote operations.

Amazon Quick supports two authentication methods for Microsoft OneNote. Choose the method that best fits your organization's security requirements.

- **Default OAuth app** – Uses an AWS-managed OAuth application. No additional credentials are needed. Users authenticate directly with their Microsoft 365 account.
- **Custom OAuth app** – Uses a customer-managed application registered in Microsoft Entra ID. This option gives your organization full control over the OAuth configuration.

For more information about the authentication methods that Amazon Quick supports, see [the section called “Authentication methods”](#).

Before you begin

Make sure that you have the following before you set up the integration.

- An active Microsoft 365 account with access to the OneNote notebooks that you want to use.
- For **Custom OAuth app**: Access to Microsoft Entra ID (formerly Azure Active Directory) to register an application and configure Microsoft Graph permissions for OneNote.
- For Amazon Quick subscription requirements, see [the section called “Set up integrations in the console”](#).

Configuring Microsoft Entra ID

If you are using **Default OAuth app** authentication, skip this section and proceed to [the section called “Setting up the connector in Amazon Quick”](#).

For Custom OAuth app authentication, register an application in the [Microsoft Entra admin center](#) and configure Microsoft Graph delegated permissions for OneNote. Add the Amazon Quick callback URL `https://{region}.quicksight.aws.amazon.com/sn/oauthcallback` as a redirect URI. Replace `{region}` with your AWS Region (for example, `us-east-1`). For step-by-step instructions, see [Authentication and authorization basics](#) in the Microsoft Graph documentation. Record the Application (client) ID and a client secret value — you need them when you configure Amazon Quick.

Add the following as Delegated permissions in your Entra app registration. For the full permissions reference, see [Microsoft Graph permissions reference](#) in the Microsoft documentation.

OneNote action integration – delegated permissions

Permission	Description
<code>Notes.Create</code>	Allows the app to create OneNote notebooks on behalf of the signed-in user.
<code>Notes.Read</code>	Allows the app to read OneNote notebooks on behalf of the signed-in user.
<code>Notes.Read.All</code>	Allows the app to read all OneNote notebooks that the signed-in user has access to.
<code>Notes.ReadWrite</code>	Allows the app to read and write OneNote notebooks on behalf of the signed-in user.
<code>Notes.ReadWrite.All</code>	Allows the app to read and write all OneNote notebooks that the signed-in user has access to.
<code>Notes.ReadWrite.CreatedByApp</code>	Allows the app to read and write OneNote notebooks created by the app on behalf of the signed-in user.
<code>offline_access</code>	Allows the app to refresh access tokens without requiring the user to sign in again. This reduces how often users need to re-authenticate.
<code>Sites.Read.All</code>	Allows the app to read documents and list items in all site collections on behalf of the signed-in user.
<code>Team.ReadBasic.All</code>	Allows the app to read the names and descriptions of teams on behalf of the signed-in user.
<code>User.Read</code>	Allows the app to read the signed-in user's profile.
<code>User.ReadBasic.All</code>	Allows the app to read a basic set of profile properties of all users on behalf of the signed-in user.

Setting up the connector in Amazon Quick

Connect from the Available tab

If you want to use Default OAuth app authentication, you can connect directly from the **Available** tab without additional configuration.

1. In the Amazon Quick console, choose **Connectors**.
2. On the **Available** tab, find **OneNote** and choose **Connect**.
3. Complete the Microsoft sign-in flow and grant the requested permissions.

To configure a connector with Custom OAuth app instead, use the **Create for your team** tab as described below.

Create from the Create for your team tab

1. In the Amazon Quick console, choose **Connectors**.
2. Choose the **Create for your team** tab.
3. Find and choose **OneNote**.
4. Enter a **Name** for your connector. Optionally, choose **+ Add Description** to add a description.
5. For **Connection type**, choose **Public network**.
6. For **OAuth Configuration**, choose one of the following authentication methods and configure the required fields.

- a. For **Default OAuth app**:

No additional credentials are needed. Choose **Next** to continue.

- b. For **Custom OAuth app**, configure the following fields:

- **Base URL** (Optional) – The Microsoft Graph base URL. Example: `https://graph.microsoft.com/v1.0`
- **Client ID** – The Application (client) ID from your Microsoft Entra ID app registration.
- **Public OAuth client** (Optional) – Select this option if your Microsoft Entra ID app is configured as a public client (no client secret).
- **Client secret** – The client secret value from your Microsoft Entra ID app registration.
- **Token URL** – The Microsoft identity platform token endpoint. Example: `https://login.microsoftonline.com/common/oauth2/v2.0/token`

- **Authorization URL** – The Microsoft identity platform authorization endpoint. Example: `https://login.microsoftonline.com/common/oauth2/v2.0/authorize`
- **Redirect URL** – Pre-filled with the Amazon Quick callback URL.

 Note

For tenant-specific authentication, replace common in the Token URL and Authorization URL with your Microsoft Entra ID tenant ID.

7. Choose **Next**.
8. A Microsoft sign-in window opens. Review the requested permissions and choose **Accept**.
9. On the **Review** page, review the available actions for the connector. Choose **Next**.
10. On the **Publish** page, choose who can access the connector. You can enable access for everyone in your organization or search for specific teams or groups.
11. Choose **Publish**.

Available actions

After you set up the connector, the actions exposed by OneNote are available. To see the current set of actions for your connector, go to the connector's **Available actions** view in the Amazon Quick console.

Managing and troubleshooting

To edit, share, or delete your connector, see [the section called “Managing existing integrations”](#).

Authentication issues

- **Sign-in fails (Default OAuth app or Custom OAuth app)** – Verify that your Microsoft 365 account is active and that you can sign in to OneNote on the web directly. For Custom OAuth app, confirm that the redirect URI in your Microsoft Entra ID app matches the Amazon Quick callback URL.
- **Insufficient permissions** – Verify that the Microsoft Graph permissions configured on your Microsoft Entra ID app include the OneNote scopes required for the operations you want to use, and that admin consent has been granted if required.

- **Invalid client credentials (Custom OAuth app)** – Verify that the Client ID and Client secret match the values in your Microsoft Entra ID app registration. Confirm that the client secret has not expired.

Microsoft Outlook integration

With the Microsoft Outlook connector, you can access Outlook's email, calendar, and contact APIs directly in Amazon Quick through natural language.

Amazon Quick supports multiple authentication methods for Microsoft Outlook. Choose the method that best fits your organization's security requirements.

- **Default OAuth app** – Uses an AWS-managed OAuth application. No additional credentials are needed. Users authenticate directly with their Microsoft account.
- **Custom OAuth app** – Uses a customer-managed application registered in Microsoft Entra. This option gives your organization full control over the OAuth configuration. Users authenticate on behalf of a signed-in user (delegated permissions).
- **Service-to-Service OAuth** – Uses client credentials for server-to-server authentication without user interaction (application permissions). Suitable for automated workflows.

For more information about the authentication methods that Amazon Quick supports, see [the section called "Authentication methods"](#).

Before you begin

Make sure that you have the following before you set up the integration.

- A Microsoft 365 account with Outlook or Exchange Online access.
- For **Custom OAuth app** or **Service-to-Service OAuth**: Access to the [Microsoft Entra admin center](#) on the Microsoft website with at least Application Developer permissions.
- For Amazon Quick subscription requirements, see [the section called "Set up integrations in the console"](#).

Configuring Microsoft Entra

If you are using **Default OAuth app** authentication, skip this section and proceed to [the section called "Setting up the connector in Amazon Quick"](#).

Before you configure Amazon Quick, create an app registration in Microsoft Entra. Complete all of the following steps in Entra before moving to the Amazon Quick console.

For more information about app registrations, see [Register an application with the Microsoft identity platform](#) in the Microsoft documentation.

Register the application

1. Open the [Microsoft Entra admin center](#).
2. In the left navigation, choose **Entra ID**, then choose **App registrations**.
3. Choose **New registration**.
4. For **Name**, enter a descriptive name for your integration.
5. For **Supported account types**, choose **Accounts in this organizational directory only**.
6. For **Redirect URI**, select **Web** and enter `https://{region}.quicksight.aws.amazon.com/sn/oauthcallback`. Replace `{region}` with the AWS Region where your Amazon Quick instance is deployed.
7. Choose **Register**.
8. On the overview page, copy the **Application (client) ID** and **Directory (tenant) ID**. You need these values for the Amazon Quick configuration.

Create a client secret

Amazon Quick needs a client secret to authenticate with Microsoft Entra. This secret acts as a password for the app registration.

1. From your app registration, choose **Certificates & secrets**.
2. Choose **New client secret**.
3. Enter a description and choose an expiration period.
4. Choose **Add**.
5. Copy the **Value** immediately. This value is only displayed once.

Important

Copy the secret **Value**, not the Secret ID. The Value is the longer string used for authentication.

Configure API permissions

Microsoft Graph supports two permission types for this integration. Delegated permissions allow the app to act on behalf of a signed-in user. Application permissions allow the app to act without a signed-in user. For more information, see [Overview of Microsoft Graph permissions](#) in the Microsoft documentation.

1. From your app registration, choose **API permissions**.
2. Choose **Add a permission**, then choose **Microsoft Graph**.
3. Choose **Delegated permissions** or **Application permissions** based on your authentication method, and add the permissions from the appropriate table below.
4. Choose **Grant admin consent for [your tenant name]** to approve the permissions.

For user authentication (delegated permissions):

Add the following as Delegated permissions in your Entra app registration. For the full permissions reference, see [Microsoft Graph permissions reference](#) in the Microsoft documentation.

Outlook action integration – delegated permissions

Permission	Description
Mail.ReadWrite	Allows the app to create, read, update, and delete email in user mailboxes.
Mail.Send	Allows the app to send mail as users in the organization.
Mail.Read.Shared	Allows the app to read mail that the user can access, including the user's own mail and shared mail.
Calendars.ReadWrite	Allows the app to create, read, update, and delete events in user calendars.
Calendars.ReadWrite.Shared	Allows the app to create, read, update and delete events in all calendars the user has permissions to access, including delegate and shared calendars.
User.Read	Allows users to sign in to the app and allows the app to read the profile of signed-in users.

Permission	Description
User.Read.All	Allows the app to read the full set of profile properties of other users in your organization.
Contacts.Read	Allows the app to read user contacts.
Place.Read.All	Allows the app to read company places (conference rooms and room lists) for calendar events and other applications.
MailboxSettings.Read	Allows the app to read the user's mailbox settings.
offline_access	Allows the app to refresh access tokens without requiring the user to sign in again. This reduces how often users need to re-authenticate.

Note

User.Read.All and Place.Read.All require administrator consent. An administrator must grant consent before users can authenticate.

For service authentication (application permissions):

Add the following as Application permissions in your Entra app registration.

Outlook action integration – application permissions

Permission	Description
Mail.ReadWrite	Allows the app to create, read, update, and delete mail in all mailboxes.
Mail.Send	Allows the app to send mail as any user.
Calendars.ReadWrite	Allows the app to create, read, update, and delete events of all calendars.
User.Read.All	Allows the app to read user profiles.

Permission	Description
Contacts.Read	Allows the app to read all contacts in all mailboxes.
Place.Read.All	Allows the app to read company places (conference rooms and room lists) for calendar events and other applications.
MailboxSettings.Read	Allows the app to read user's mailbox settings.

Important

With service authentication, all actions execute as the service account. Any user with access to this integration can perform actions across all mailboxes that the service account can access. Scope the application permissions appropriately for your organization's security requirements.

Record your credentials

Before leaving the Microsoft Entra admin center, confirm you have the following values. You need them for the Amazon Quick configuration.

Required credentials from Microsoft Entra

Value	Where to find it
Application (client) ID	App registration overview page
Directory (tenant) ID	App registration overview page
Client secret value	Certificates & secrets page

Setting up the connector in Amazon Quick

Connect from the Available tab

If you want to use Default OAuth app authentication, you can connect directly from the **Available** tab without additional configuration.

1. In the Amazon Quick console, choose **Connectors**.
2. On the **Available** tab, find **Microsoft Outlook** and choose **Connect**.
3. Complete the Microsoft sign-in flow and grant the requested permissions.

To configure a connector with Custom OAuth app or Service-to-Service OAuth instead, use the **Create for your team** tab as described below.

Create from the Create for your team tab

After you complete any required Entra configuration, create the connector in Amazon Quick.

1. In the Amazon Quick console, choose **Connectors**.
2. Choose the **Create for your team** tab.
3. Find and choose **Microsoft Outlook**.

Note

If a Microsoft Outlook connector already exists, a dialog appears with your existing connectors. To use an existing connector, choose it. To create a new one, choose **No, create new**.

4. Enter a **Name** for your connector. Optionally, choose **+ Add Description** to add a description.
5. For **Connection type**, choose **Public network**.
6. For **OAuth Configuration**, choose one of the following authentication methods and configure the required fields.

a. For **Default OAuth app**:

No additional credentials are needed. Choose **Next** to continue.

b. For **Custom OAuth app** (user authentication with delegated permissions), configure the following fields:

- **Base URL** (Optional) – The Microsoft Graph API base URL. Example: `https://graph.microsoft.com/v1.0`
- **Client ID** – The Application (client) ID from your Entra app registration.
- **Client secret** – The client secret value from your Entra app registration.

- **Token URL** – The token endpoint. Example: `https://login.microsoftonline.com/{tenant-id}/oauth2/v2.0/token`
 - **Authorization URL** – The authorization endpoint. Example: `https://login.microsoftonline.com/{tenant-id}/oauth2/v2.0/authorize`
 - **Redirect URL** – Pre-filled with the Amazon Quick callback URL.
- c. For **Service-to-Service OAuth** (service authentication with application permissions), configure the following fields:
- **Base URL** (Optional) – The Microsoft Graph API base URL. Example: `https://graph.microsoft.com/v1.0`
 - **Client ID** – The Application (client) ID from your Entra app registration.
 - **Client secret** – The client secret value from your Entra app registration.
 - **Token URL** – The token endpoint. Example: `https://login.microsoftonline.com/{tenant-id}/oauth2/v2.0/token`

 Note

The scope for the client credentials token request (`https://graph.microsoft.com/.default`) is set automatically by Amazon Quick. You do not need to configure it manually.

7. Choose **Next**.
8. If you chose **Default OAuth app** or **Custom OAuth app**, a Microsoft authorization window opens. Review the requested permissions and choose **Accept**.

If you see an error instead of the consent dialog, your organization might restrict third-party app access. See [the section called “Admin consent for Microsoft 365”](#).

9. On the **Review** page, review the available actions for the connector. Choose **Next**.
10. On the **Publish** page, choose who can access the connector. You can enable access for everyone in your organization or search for specific teams or groups.
11. Choose **Publish**.

Available actions

After you set up the connector, the following actions are available.

Microsoft Outlook available actions

Category	Action	Description
Email	List User Mails	Lists emails in a mailbox.
Email	List Folder Messages	Lists messages in a specific mail folder.
Email	View Email	Retrieves email details by ID.
Email	Send User Email	Sends a new email message.
Email	Reply To Email	Replies to an existing email.
Email	Forward User Email	Forwards an email to other recipients.
Email	Update Email	Edits email properties.
Email	Delete Email	Removes an email from a mailbox.
Email	Move Email To Folder	Moves an email to a different folder.
Email	List Email Attachments	Lists attachments on an email.
Email	Get Attachment	Retrieves attachment details and content by ID.
Calendar	List Calendar Events	Lists events on a calendar.
Calendar	List Calendar View	Lists meetings in a specified date range.
Calendar	Create Calendar Event	Creates a new meeting or appointment.
Calendar	Update Calendar Event	Modifies an existing event.
Calendar	Delete Calendar Event	Removes an event from a calendar.
Calendar	Find Meeting Times	Suggests meeting times based on attendee availability.
Contacts	List Contacts	Lists contacts.

Category	Action	Description
Users	List Users	Lists users in the organization.
Settings	Get Mailbox Settings	Reads mailbox configuration.
Places	List Places	Lists meeting rooms and room lists.

Manage and troubleshoot

To edit, share, or delete your integration, see [the section called “Managing existing integrations”](#).

Authentication issues

- **Incorrect app registration** – Verify the app registration in Microsoft Entra includes the required API permissions and that admin consent has been granted.
- **Expired client secret** – Check if the client secret has expired in **Certificates & secrets** and generate a new one if needed.
- **Incorrect redirect URI** – Verify the redirect URI in Microsoft Entra matches `https://{region}.quicksight.aws.amazon.com/sn/oauthcallback`.

Common error messages

- **Access denied. You do not have permission to perform this action** – The authenticated user does not have the required permissions. Contact your administrator to verify and grant appropriate permissions.
- **AADSTS50020: User account from identity provider does not exist in tenant** – The user account is not configured in the correct Microsoft Entra tenant. Verify the user account exists in the tenant that matches the Directory (tenant) ID in your app registration.

Admin consent for Microsoft 365

When you use the **Default OAuth app** authentication method, Amazon Quick uses an AWS-managed application to access Microsoft Outlook on behalf of the signed-in user. Most users can complete setup without any extra steps. However, if your Microsoft 365 tenant restricts third-party app access, a Microsoft 365 administrator must grant one-time consent before users can connect.

If you see an error when you sign in during connector setup, your organization might restrict third-party app access. Share the following information with your Microsoft 365 administrator:

- **What to do:** Grant admin consent for the Amazon Quick Microsoft Outlook integration application.
- **Why:** Amazon Quick needs delegated access to Outlook email, calendar, and contact data to perform actions on behalf of users.

An administrator can grant consent in one of the following ways:

- **Through the consent dialog** – A Global Administrator or Privileged Role Administrator initiates the connector setup flow. In the Microsoft sign-in dialog, they select the **Consent on behalf of your organization** check box and choose **Accept**.
- **Through the Microsoft Entra admin center** – Sign in to the [Microsoft Entra admin center](#) on the Microsoft website. Choose **Enterprise applications**, locate the Amazon Quick application, choose **Permissions**, and choose **Grant admin consent for *Your Organization***.

After consent is granted, any user in your organization can connect without being prompted for individual consent.

Note

To check whether your tenant restricts user consent, go to the Microsoft Entra admin center and choose **Enterprise applications, Consent and permissions, User consent settings**. If the setting is **Do not allow user consent**, an administrator must grant consent before users can use the connector.

Microsoft SharePoint integration

With the Microsoft SharePoint integration, you can manage SharePoint lists, items, and files. You can also create knowledge bases from SharePoint content for AI-powered search and Q&A.

What you can do

Connector

Manage SharePoint lists, items, files, and Excel workbooks through Microsoft Graph API calls.

Knowledge base

Index SharePoint document libraries, sites, and pages. Amazon Quick agents can then search and answer questions about the content. Amazon Quick supports two setup methods: user-managed setup for quick configuration, and admin-managed setup with service credentials for document-level access control.

Next steps

After you review the prerequisites and capabilities, continue with one of the following topics:

- To set up the connector, see [the section called “Action integration”](#).
- To create a knowledge base from SharePoint content, see [the section called “Knowledge base integration”](#).

Microsoft SharePoint action integration

Use the Microsoft SharePoint connector to manage lists, items, files, and Excel workbooks directly in Amazon Quick through natural language.

Amazon Quick supports multiple authentication methods for Microsoft SharePoint. Choose the method that best fits your organization's security requirements.

- **Default OAuth app** – Uses an AWS-managed OAuth application. No additional credentials are needed. Users authenticate directly with their Microsoft account.
- **Custom OAuth app** – Uses a customer-managed application registered in Microsoft Entra. This option gives your organization full control over the OAuth configuration. Users authenticate on behalf of a signed-in user (delegated permissions).
- **Service-to-Service OAuth** – Uses client credentials for server-to-server authentication without user interaction (application permissions). Suitable for automated workflows.

For more information about the authentication methods that Amazon Quick supports, see [the section called “Authentication methods”](#).

Before you begin

Make sure you have the following before you set up the integration.

- A Microsoft 365 account with SharePoint access.
- For **Custom OAuth app** or **Service-to-Service OAuth**: Access to the [Microsoft Entra admin center](#) on the Microsoft website with at least Application Developer permissions.
- For subscription requirements, see [the section called “Set up integrations in the console”](#).

Configure Microsoft Entra

If you are using **Default OAuth app** authentication, skip this section and proceed to [the section called “Setting up the connector in Amazon Quick”](#).

Before you configure Amazon Quick, create an app registration in Microsoft Entra. Complete all of the following steps in Entra before moving to the Amazon Quick console.

For more information about app registrations, see [Register an application with the Microsoft identity platform](#) in the Microsoft documentation.

Register the application

1. Open the [Microsoft Entra admin center](#).
2. In the left navigation, choose **Entra ID**, then choose **App registrations**.
3. Choose **New registration**.
4. For **Name**, enter a descriptive name for your integration.
5. For **Supported account types**, choose **Accounts in this organizational directory only**.
6. For **Redirect URI**, select **Web** and enter `https://{region}.quicksight.aws.amazon.com/sn/oauthcallback`. Replace `{region}` with the AWS Region where your Amazon Quick instance is deployed.
7. Choose **Register**.
8. On the overview page, copy the **Application (client) ID** and **Directory (tenant) ID**. You need these values for the Amazon Quick configuration.

Create a client secret

Amazon Quick needs a client secret to authenticate with Microsoft Entra. This secret acts as a password for the app registration.

1. From your app registration, choose **Certificates & secrets**.
2. Choose **New client secret**.

3. Enter a description and choose an expiration period.
4. Choose **Add**.
5. Copy the **Value** immediately. This value is only displayed once.

Important

Copy the secret **Value**, not the Secret ID. The Value is the longer string used for authentication.

Configure API permissions

Microsoft Graph supports two permission types for this integration. Delegated permissions allow the app to act on behalf of a signed-in user. Application permissions allow the app to act without a signed-in user. For more information, see [Overview of Microsoft Graph permissions](#) in the Microsoft documentation.

1. From your app registration, choose **API permissions**.
2. Choose **Add a permission**, then choose **Microsoft Graph**.
3. Choose **Delegated permissions** or **Application permissions** based on your authentication method, and add the permissions from the appropriate table below.
4. Choose **Grant admin consent for [your tenant name]** to approve the permissions.

For user authentication (delegated permissions):

Add the following as Delegated permissions in your Entra app registration. For the full permissions reference, see [Microsoft Graph permissions reference](#) in the Microsoft documentation.

SharePoint action integration – delegated permissions

Permission	Description
Files.ReadWrite	Allows the app to read, create, update, and delete the signed-in user's files.
Sites.ReadWrite.All	Allows the application to edit or delete documents and list items in all site collections on behalf of the signed-in user.

Permission	Description
User.Read	Allows the app to read the signed-in user's profile.
offline_access	Allows the app to refresh access tokens without requiring the user to sign in again. This reduces how often users need to re-authenticate.

For service authentication (application permissions):

Add the following as Application permissions in your Entra app registration.

SharePoint action integration – application permissions

Permission	Description
Sites.ReadWrite.All	Allows the app to create, read, update, and delete documents and list items in all site collections without a signed-in user.

Important

With service authentication, all actions execute as the service account. Any user with access to this integration can perform actions across all site collections that the service account can access. Scope the application permissions appropriately for your organization's security requirements.

Record your credentials

Before leaving the Microsoft Entra admin center, confirm you have the following values. You need them for the Amazon Quick configuration.

Required credentials from Microsoft Entra

Value	Where to find it
Application (client) ID	App registration overview page
Directory (tenant) ID	App registration overview page

Value	Where to find it
Client secret value	Certificates & secrets page

Setting up the connector in Amazon Quick

Connect from the Available tab

If you want to use Default OAuth app authentication, you can connect directly from the **Available** tab without additional configuration.

1. In the Amazon Quick console, choose **Connectors**.
2. On the **Available** tab, find **SharePoint** and choose **Connect**.
3. Complete the Microsoft sign-in flow and grant the requested permissions.

To configure a connector with Custom OAuth app or Service-to-Service OAuth instead, use the **Create for your team** tab as described below.

Create from the Create for your team tab

After you complete any required Entra configuration, create the connector in Amazon Quick.

1. In the Amazon Quick console, choose **Connectors**.
2. Choose the **Create for your team** tab.
3. Find and choose **Microsoft SharePoint**.

Note

If a Microsoft SharePoint connector already exists, a dialog appears with your existing connectors. To use an existing connector, choose it. To create a new one, choose **No, create new**.

4. On the **Integration type** page, select **Perform actions in Microsoft SharePoint Online** and choose **Next**.
5. Enter a **Name** for your connector. Optionally, choose **+ Add Description** to add a description.
6. For **Connection type**, choose **Public network**.

7. For **OAuth Configuration**, choose one of the following authentication methods and configure the required fields.

a. For **Default OAuth app**:

No additional credentials are needed. Choose **Next** to continue.

b. For **Custom OAuth app** (user authentication with delegated permissions), configure the following fields:

- **Base URL** (Optional) – The Microsoft Graph API base URL. Example: `https://graph.microsoft.com/v1.0`
- **Client ID** – The Application (client) ID from your Entra app registration.
- **Client secret** – The client secret value from your Entra app registration.
- **Token URL** – The token endpoint. Example: `https://login.microsoftonline.com/{tenant-id}/oauth2/v2.0/token`
- **Authorization URL** – The authorization endpoint. Example: `https://login.microsoftonline.com/{tenant-id}/oauth2/v2.0/authorize`
- **Redirect URL** – Pre-filled with the Amazon Quick callback URL.

c. For **Service-to-Service OAuth** (service authentication with application permissions), configure the following fields:

- **Base URL** (Optional) – The Microsoft Graph API base URL. Example: `https://graph.microsoft.com/v1.0`
- **Client ID** – The Application (client) ID from your Entra app registration.
- **Client secret** – The client secret value from your Entra app registration.
- **Token URL** – The token endpoint. Example: `https://login.microsoftonline.com/{tenant-id}/oauth2/v2.0/token`

 Note

The scope for the client credentials token request (`https://graph.microsoft.com/.default`) is set automatically by Amazon Quick. You do not need to configure it manually.

8. Choose **Next**.

9. If you chose **Default OAuth app** or **Custom OAuth app**, a Microsoft authorization window opens. Review the requested permissions and choose **Accept**.

If you see an error instead of the consent dialog, your organization might restrict third-party app access. See [the section called "Admin consent for Microsoft 365"](#).

10. On the **Review** page, review the available actions for the connector. Choose **Next**.
11. On the **Publish** page, choose who can access the connector. You can enable access for everyone in your organization or search for specific teams or groups.
12. Choose **Publish**.

Available actions

After you set up the integration, the following actions are available.

Microsoft SharePoint available actions

Category	Action	Description
Lists and items	View items	Get the collection of items in a list.
Lists and items	Get Item	Returns the metadata for an item in a list.
Lists and items	Get List	Returns the metadata for a list.
Lists and items	Update Item	Update the properties on a list item.
Lists and items	Delete Item	Removes an item from a list.
Files	Upload File	Upload a new file or update an existing file. Supports files up to 250 MB.
Files	Search Site Drive Items	Search the hierarchy of items matching a query.

Category	Action	Description
Excel workbooks	List Sheets	Retrieve a list of worksheet objects.
Excel workbooks	Add Sheet	Add a new worksheet to the workbook.
Excel workbooks	Read Sheet	Retrieve the properties of a worksheet object.
Excel workbooks	Update Sheet	Update the properties of a worksheet object.
Excel workbooks	Delete Sheet	Delete the worksheet from the workbook.
Excel workbooks	Read Cell	Get the value of a single cell by row and column number.
Excel workbooks	Write Cell	Set the value of a single cell by row and column number.
Excel workbooks	Read Range	Get the values of a range.
Excel workbooks	Write Range	Update the values of a range.
Excel workbooks	Clear Range	Clear range values, format, fill, and border.
Excel workbooks	Delete Range	Delete the cells associated with the range.
Excel workbooks	Get Used Range	Get the smallest range that encompasses cells with a value or formatting.

Manage and troubleshoot

To edit, share, or delete your integration, see [the section called “Managing existing integrations”](#).

Authentication issues

- **Incorrect app registration** – Verify the app registration in Microsoft Entra includes the required API permissions and that admin consent has been granted.
- **Expired client secret** – Check if the client secret has expired in **Certificates & secrets** and generate a new one if needed.
- **Incorrect redirect URI** – Verify the redirect URI in Microsoft Entra matches `https://{region}.quicksight.aws.amazon.com/sn/oauthcallback`.

Common error messages

- **Access denied. You do not have permission to perform this action** – The authenticated user does not have the required permissions. Contact your administrator to verify and grant appropriate permissions.
- **AADSTS50020: User account from identity provider does not exist in tenant** – The user account is not configured in the correct Microsoft Entra tenant. Verify the user account exists in the tenant that matches the Directory (tenant) ID in your app registration.

Admin consent for Microsoft 365

When you use the **Default OAuth app** authentication method, Amazon Quick uses an AWS-managed application to access Microsoft SharePoint on behalf of the signed-in user. Most users can complete setup without any extra steps. However, if your Microsoft 365 tenant restricts third-party app access, a Microsoft 365 administrator must grant one-time consent before users can connect.

If you see an error when you sign in during connector setup, your organization might restrict third-party app access. Share the following information with your Microsoft 365 administrator:

- **What to do:** Grant admin consent for the Amazon Quick Microsoft SharePoint integration application.
- **Why:** Amazon Quick needs delegated access to SharePoint sites, lists, files, and Excel workbooks to perform actions on behalf of users.

An administrator can grant consent in one of the following ways:

- **Through the consent dialog** – A Global Administrator or Privileged Role Administrator initiates the connector setup flow. In the Microsoft sign-in dialog, they select the **Consent on behalf of your organization** check box and choose **Accept**.
- **Through the Microsoft Entra admin center** – Sign in to the [Microsoft Entra admin center](#) on the Microsoft website. Choose **Enterprise applications**, locate the Amazon Quick application, choose **Permissions**, and choose **Grant admin consent for *Your Organization***.

After consent is granted, any user in your organization can connect without being prompted for individual consent.

 Note

To check whether your tenant restricts user consent, go to the Microsoft Entra admin center and choose **Enterprise applications, Consent and permissions, User consent settings**. If the setting is **Do not allow user consent**, an administrator must grant consent before users can use the connector.

Microsoft SharePoint knowledge base integration

Microsoft SharePoint Online is a collaborative platform for building websites, managing document libraries, and organizing content across your organization. With the Amazon Quick SharePoint connector, you can create knowledge bases from your SharePoint content to power AI-driven search and Q&A.

Amazon Quick supports two authentication methods for connecting to SharePoint Online:

- **User-managed setup** – You sign in to SharePoint directly to authorize the connection. This is the simplest way to get started. For more information, see [the section called “User-managed setup”](#).
- **Admin-managed setup (service credentials)** – An Entra ID app registration authenticates on behalf of the application using certificate-based credentials. A key benefit of admin-managed setup is built-in document-level access control (ACL). Amazon Quick automatically syncs access control lists from SharePoint and verifies each user's permissions at query time. Users see answers only from documents that they are authorized to access. For more information, see [the section called “Admin-managed setup”](#).

After connecting, Amazon Quick indexes your SharePoint document libraries, sites, lists, and pages into a knowledge base. Your Amazon Quick agents can then search this content and generate answers grounded in your SharePoint data.

Prerequisites

Before you set up the SharePoint knowledge base integration, make sure that you have the following:

- An AWS account with an active Amazon Quick instance.
- A Microsoft 365 account with SharePoint Online access.
- For subscription requirements, see [the section called “Set up integrations in the console”](#).

For admin-managed setup, additional prerequisites apply. For more information, see [the section called “Admin-managed setup”](#).

Tip

Custom SharePoint domains are supported. If your organization uses a custom domain instead of the default `https://tenant.sharepoint.com` format, you can use your custom SharePoint URL when configuring the knowledge base.

User-managed setup

With user-managed setup, you sign in to SharePoint directly and Amazon Quick handles authentication. Most users can complete setup in a few minutes.

Prerequisites

Before you begin, make sure you have the following:

- A Microsoft 365 account with access to the SharePoint sites you want to index.
- Access to Amazon Quick with permissions to create knowledge bases.
- A browser that allows popups from the Amazon Quick console domain.

Important

User-managed setup does not support document-level access control (ACL). All indexed content is accessible to any user who has access to the knowledge base in Amazon Quick. Individual permissions in SharePoint are not enforced. Carefully review which content you include when creating a knowledge base. If you require document-level access control, use [Admin-managed setup \(service credentials\)](#) instead.

Create the SharePoint knowledge base

Navigate to Knowledge

1. In Amazon Quick, choose **Knowledge** from the left navigation panel.
2. Under **Set up new knowledge base**, locate **Microsoft SharePoint Online** and choose the **+** icon.

Sign in to SharePoint

The **Create SharePoint knowledge base** wizard opens on the Authentication method step. **Quick setup** is selected by default.

1. Under **Quick setup**, choose **Sign in to SharePoint**.
2. A Microsoft sign-in window opens. Enter your Microsoft 365 credentials.
3. If a permissions consent dialog appears, review the permissions and choose **Accept**.

If you see an error instead of the consent dialog, your organization might restrict third-party app access. See [the section called "Admin consent for Microsoft 365"](#).

4. After successful authentication, choose **Next**.

Choose content

1. Enter a **Name** and optional **Description** for your knowledge base.
2. In the **Content** section, choose **Add content**.
3. A dialog opens showing SharePoint sites and content accessible to your account. Select the pages, lists, files, or folders you want to index.

4. Choose **Add** to confirm your selections. You can add more content or remove items before continuing.
5. Choose **Create**.

Tip

If the site you are looking for does not appear in the dialog, choose the search link at the top of the Content section. Enter the full URL of the SharePoint site (for example, `https://contoso.sharepoint.com/sites/marketing`) and choose **Browse**. The dialog then displays the document libraries, lists, pages, folders, and files within that site. Selecting a folder includes all files and subfolders within it.

Initial sync

After you choose **Create**, you are returned to the knowledge base list page. The knowledge base might take a few minutes to finish provisioning. Once creation is complete, an initial sync is automatically triggered.

Admin consent for Microsoft 365

Most users complete setup without any extra steps. However, if your Microsoft 365 tenant restricts third-party app access, you might see an error when you sign in. In this case, a Microsoft 365 administrator needs to grant one-time consent for the Amazon Quick application. After consent is granted, any user in your organization can connect.

If you are not a Microsoft 365 administrator, share the following information with your administrator:

- **What to do:** Grant admin consent for the Amazon Quick SharePoint integration application.
- **Why:** Amazon Quick needs delegated read access to SharePoint sites and files to index content for knowledge bases.

Grant organization-wide admin consent

Some Amazon Quick features require delegated permissions from Microsoft Entra. By default, each user sees a Microsoft consent dialog the first time they use the feature. A Microsoft 365

administrator can pre-consent on behalf of the entire organization so that individual users aren't prompted. This is a one-time action per application.

 Note

If your Microsoft 365 tenant is configured to restrict user consent for third-party applications, admin consent is required, not optional. Without it, users see an error when they attempt to use the feature.

The following table describes the user experience with and without admin consent.

Admin consent scenarios

Scenario	User experience
Admin consent not granted	Each user sees the Microsoft permissions consent dialog on first use. Users might be blocked if your tenant restricts user consent for third-party apps.
Admin consent granted	Users aren't prompted for consent. The feature works immediately for all users in the organization.

Granting consent through the consent dialog

The simplest way to grant admin consent is through the Microsoft consent dialog that appears during the feature flow.

To grant consent through the consent dialog

1. Have a Global Administrator or Privileged Role Administrator initiate the feature flow that triggers the consent dialog.
2. In the Microsoft sign-in dialog, select the **Consent on behalf of your organization** check box.
3. Choose **Accept**.

This grants consent for the requested delegated permissions for all users in your Microsoft 365 tenant.

Granting consent through the Microsoft Entra admin center

Administrators can also grant consent directly from the Microsoft Entra admin center.

To grant consent through the Microsoft Entra admin center

1. Sign in to the [Microsoft Entra admin center](#).
2. In the left navigation pane, expand **Entra ID** and choose **Enterprise applications**.
3. Locate the enterprise application for the Amazon Quick feature.

Note

The application name appears in the consent dialog that users see when they first use the feature.

4. In the left navigation pane, choose **Permissions**.
5. Choose **Grant admin consent for *Your Organization***.
6. Confirm the consent.

Verifying consent

After you grant consent, the Enterprise application's **Permissions** page shows all delegated permissions with a status indicator under the **Admin consent** column.

Note

When an administrator grants organizational consent, Microsoft Entra automatically creates an Enterprise Application (service principal) in your tenant. To revoke access, disable or delete this service principal from **Enterprise applications** in the Microsoft Entra admin center.

Checking tenant consent settings

To check whether your tenant restricts user consent, complete the following steps.

To check tenant consent settings

1. In the Microsoft Entra admin center, choose **Entra ID, Enterprise applications, Consent and permissions, User consent settings**.
2. If the setting is **Do not allow user consent**, an administrator must grant consent before users can use the feature.

Permissions requested

The following delegated permissions are requested when a user signs in. Share this list with your administrator if they need to review the permissions before granting consent.

User-managed setup – permissions

Permission	API	Type	Description
Files.Read.All	Microsoft Graph	Delegated	Read all files the signed-in user can access.
Notes.Read.All	Microsoft Graph	Delegated	Read all OneNote notebooks the signed-in user can access.
User.Read	Microsoft Graph	Delegated	Sign in and read the user's profile.
Sites.Read.All	Microsoft Graph	Delegated	Read documents and list items in all site collections.
offline_access	Microsoft Graph	Delegated	Maintain access using refresh tokens.
AllSites.Read	Office 365 SharePoint Online	Delegated	Read items in all site collections.

Manage and troubleshoot user-managed connections

To edit, share, or delete your integration, see [the section called “Managing existing integrations”](#).

- **App blocked by administrator** – Your Microsoft 365 tenant restricts third-party app access. Ask your Microsoft 365 administrator to grant admin consent. For more information, see [the section called “Admin consent for Microsoft 365”](#).
- **Sign-in window closes without completing** – Verify that your browser allows popups from the Amazon Quick console domain and that third-party cookies are enabled.
- **Token expired** – Delegated credentials last approximately 90 days. If syncs fail after this period, you need to re-authenticate. For more information, see [the section called “Token refresh with user-managed setup”](#).
- **Missing content** – Verify that the account you used for authentication has access to the files and folders you selected. Content shared with you after the initial sync requires a resync to be indexed.

For additional troubleshooting, including sync monitoring and reports, see [the section called “Troubleshooting”](#).

Next steps

After your knowledge base is created and the initial sync completes, you can use it in Amazon Quick to answer questions from your SharePoint content. To manage your knowledge base, see [the section called “Managing existing integrations”](#).

Admin-managed setup (service credentials)

With admin-managed setup, an administrator configures an AWS KMS signing key and a Microsoft Entra ID app registration with certificate-based credentials. Individual users don't need to authorize through sign-in.

Admin-managed setup optionally includes document-level access control list (ACL) support. When enabled, Amazon Quick syncs ACLs from SharePoint and verifies each user's permissions at query time. For more information, see [the section called “Document-level access controls”](#).

Prerequisites

Before you begin, make sure that you have the following:

- Administrator access to the Amazon Quick admin console.
- Administrator access to Microsoft Entra ID to register an application and grant API permissions.
- A SharePoint Online site with content to index.

Setup overview

The setup involves the following phases:

1. **Set up service credentials** – Create a KMS signing key, generate a certificate, register an application in Entra, and grant Amazon Quick permission to use the key. For more information, see [the section called “Set up service credentials”](#).
2. **Create the knowledge base in Amazon Quick** – Create a SharePoint knowledge base using the service credentials from Phase 1. For more information, see [the section called “Create the knowledge base”](#).

Document-level access control is optionally available for all admin-managed knowledge bases. For more information about how access controls work, see [the section called “Document-level access controls”](#).

Set up service credentials

Before you create the knowledge base in Amazon Quick, complete the following configuration steps in AWS and Microsoft Entra ID. You create a KMS signing key, generate a certificate, register an application in Entra, and grant Amazon Quick permission to use the key.

This setup involves multiple systems and might require coordination between different administrators in your organization. The following table summarizes each step and the role needed to complete it.

Setup steps and roles

Step	What you do	Role needed
1. KMS key	Create an asymmetric signing key in AWS KMS.	AWS administrator (KMS and IAM console access)
2. Certificate	Generate a self-signed certificate using the KMS public key.	Same as Step 1 (AWS CLI and OpenSSL required)
3. Entra app	Register an application in Microsoft Entra, assign API permissions, and upload the certificate.	Microsoft 365 Global Admin or Privileged Role Admin

Step	What you do	Role needed
3b. Sites.Selected (optional)	Create a temporary admin app and grant per-site permissions via the Microsoft Graph API.	Microsoft 365 Global Admin (same as Step 3)
4. KMS key access	Grant Amazon Quick permission to use the KMS key for signing.	Amazon Quick administrator (Admin Pro)
5. Create KB	Create the knowledge base in Amazon Quick using the credentials from the previous steps.	Any Amazon Quick user (Author Pro or Admin Pro)

 Tip

In many organizations, a single person with both AWS and Microsoft 365 administrator access can complete all steps. If responsibilities are split across teams, share this table to coordinate the setup.

Prerequisites

Before you begin, make sure that you have the following:

- An AWS account with an active Amazon Quick instance.
- Access to the AWS KMS console (for creating the signing key).
- Amazon Quick administrator access (Admin Pro role) for granting KMS key permissions.
- A Microsoft 365 tenant with SharePoint Online.
- Global Administrator or Privileged Role Administrator access in Microsoft Entra ID.
- OpenSSL 3.0 or later and AWS CLI installed locally.
- The AWS account and Amazon Quick instance must be in the same Region.

Permissions

The permissions you assign depend on two choices:

- Whether you plan to enable document-level access control (ACL crawling).

- Whether you want to grant access to all SharePoint sites or only specific sites.

Choose your permission scope

By default, the Entra app registration uses `Sites.Read.All` or `Sites.FullControl.All`, which grants access to all SharePoint sites in your tenant. If your organization requires least-privilege access, you can use `Sites.Selected` instead. With `Sites.Selected`, the app can only access sites that you explicitly grant permission to.

Permission scope comparison

Scope	Access	Additional steps
All sites (default)	App can read all SharePoint sites in the tenant.	—
<code>Sites.Selected</code>	App can only access sites you explicitly grant.	Requires a temporary admin app and a Microsoft Graph API call for each site. See the section called “Step 3b: Grant site-level permissions (Sites.Selected only)” .

Note

If you use `Sites.Selected`, you must grant access for each site individually. Any new sites added to the knowledge base in the future also require a separate permission grant.

All sites – content only (no ACL)

Content-only permissions

API	Permission	Type
Microsoft Graph	<code>Sites.Read.All</code>	Application
SharePoint REST	<code>Sites.Read.All</code>	Application

All sites – with ACL crawling

ACL crawling permissions

API	Permission	Type
Microsoft Graph	Sites.Read.All	Application
Microsoft Graph	User.Read.All	Application
Microsoft Graph	GroupMember.Read.All	Application
SharePoint REST	Sites.FullControl.All	Application

Important

Choose either the all-sites permissions in the preceding tables or the Sites.Selected permissions in the following tables. Do not combine both. If you are unsure, start with all-sites. You can create a new Entra app registration with Sites.Selected later if needed.

Sites.Selected – content only (no ACL)

Sites.Selected content-only permissions

API	Permission	Type
Microsoft Graph	Sites.Selected	Application
SharePoint REST	Sites.Selected	Application

Sites.Selected – with ACL crawling

Sites.Selected ACL crawling permissions

API	Permission	Type
Microsoft Graph	Sites.Selected	Application
Microsoft Graph	User.Read.All	Application

API	Permission	Type
Microsoft Graph	GroupMember.Read.All	Application
SharePoint REST	Sites.Selected	Application

 **Note**

OneNote crawling (Notes.Read.All) is not supported in admin-managed setup. Microsoft retired app-only tokens for OneNote APIs on March 31, 2025. Use [User-managed setup](#) for OneNote content.

Values collected during setup

The following table summarizes the values you create or collect during setup and where you use them.

Values reference

Value	Created in step	Used in step
KMS Key ARN	1 (KMS)	2 (Certificate), 4 (IAM), Quick setup
Certificate file (certificate.cer)	2 (Certificate)	3 (Entra upload)
Certificate thumbprint (base64url)	2 (Certificate)	Quick setup
Application (Client) ID	3 (Entra)	Quick setup
Directory (Tenant) ID	3 (Entra)	Quick setup
SharePoint domain URL	Your M365 tenant	Quick setup

Step 1: Create an AWS KMS asymmetric signing key

Amazon Quick uses an AWS KMS asymmetric key to sign OAuth assertions when authenticating with Microsoft Entra ID. The private key never leaves KMS. Only the public key is exported and embedded in a certificate that gets uploaded to your Entra app registration.

Create the KMS key

1. Open the [AWS KMS console](#).
2. In the left navigation, choose **Customer managed keys**.
3. Choose **Create key**.

Configure the key

On the **Configure key** page, set the following values:

KMS key configuration

Setting	Value
Key type	Asymmetric
Key usage	Sign and verify
Key spec	RSA_2048
Key material origin	KMS (recommended)
Regionality	Single-Region key (default). Multi-Region keys are not supported.

Add labels

On the **Add labels** page, enter an alias for the key. For example: quick-sharepoint-service-auth.

Note

The key administrator and key usage permissions on the following pages are optional. The defaults are sufficient for this setup. You grant Amazon Quick access to the key separately in Step 4.

Choose **Skip to review**, then choose **Finish** to create the key.

Record the Key ARN

After the key is created, open the key detail page and record the **Key ARN**. The ARN has the following format:

```
arn:aws:kms:us-west-2:123456789012:key/xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxxxx
```

You need this value in Steps 2, 4, and when you create the knowledge base in Quick.

Step 2: Generate a self-signed certificate

Microsoft Entra ID requires an X.509 certificate to validate signed assertions. Because the KMS private key never leaves AWS KMS, you cannot use it directly with OpenSSL. Instead, you generate a temporary local key pair and create a certificate signing request. Then, use the OpenSSL -force_pubkey option to inject the KMS public key into the final certificate. The result is a self-signed certificate whose public key matches the KMS key pair.

Prerequisites

- AWS CLI installed and configured.
- OpenSSL 3.0 or later.
- The KMS Key ARN from Step 1.

Generate the certificate

Run the following commands in a terminal. Replace the *placeholder* values with your own.

Verify OpenSSL version

```
openssl version
```

Confirm the output shows version 3.0 or later.

Export the KMS public key

```
aws kms get-public-key \  
  --key-id KMS_KEY_ARN \  
  --region REGION \  
  --output text \  
  --query PublicKey | base64 --decode > public_key.der
```

Note

On macOS, use `base64 --decode` or `base64 -D` depending on your shell environment.

Convert the public key to PEM format

```
openssl rsa -pubin -inform DER -in public_key.der -outform PEM -out kms_public_key.pem
```

Generate a temporary local key pair

```
openssl genrsa -out temp_private_key.pem 2048
```

Create a certificate signing request

```
openssl req -new \  
  -key temp_private_key.pem \  
  -out cert.csr \  
  -subj "/CN=QuickSharePointServiceAuth/0=YourOrganization/C=US"
```

Generate the certificate with the KMS public key

```
openssl x509 -req \  
  -in cert.csr \  
  -signkey temp_private_key.pem \  
  -out certificate.pem \  
  -days 730 \  
  -force_pubkey kms_public_key.pem
```

Note

OpenSSL displays the warning `Signature key and public key of cert do not match`. This is expected because the certificate is signed with the temporary local key but contains the KMS public key. The certificate is valid and works correctly with Microsoft Entra.

Convert to DER format for Entra upload

```
openssl x509 -in certificate.pem -outform DER -out certificate.cer
```

Clean up temporary files

```
rm -f temp_private_key.pem cert.csr public_key.der kms_public_key.pem certificate.pem
```

Important

Keep the `certificate.cer` file. You upload it to Microsoft Entra ID in Step 3.

Calculate the certificate thumbprint

Run the following command to calculate the base64url-encoded SHA-1 thumbprint of the certificate:

```
openssl dgst -sha1 -binary certificate.cer | base64 | tr '+/' '-_' | tr -d '='
```

Record this value. You enter it when you create the knowledge base in Quick.

Note

The base64url-encoded thumbprint is different from the hexadecimal thumbprint shown in the Microsoft Entra portal. Quick requires the base64url format.

Step 3: Register an application in Microsoft Entra ID

This step is required regardless of whether you use all-sites permissions or Sites.Selected. The only difference is which API permissions you assign in the [the section called "Configure API permissions"](#) section.

Register the application

1. Sign in to the [Microsoft Entra admin center](#).
2. In the left navigation, expand **Entra ID** and choose **App registrations**.
3. Choose **New registration**.
4. For **Name**, enter QuickSharePointServiceAuth.
5. For **Supported account types**, select **Accounts in this organizational directory only (Single tenant)**.
6. Leave **Redirect URI** blank. A redirect URI is not required because the application uses the client credentials flow, not an interactive sign-in flow.
7. Choose **Register**.

Record application details

On the application **Overview** page, record the following values:

Application details

Value	Location
Application (client) ID	Shown on the Overview page under Essentials .
Directory (tenant) ID	Shown on the Overview page under Essentials .

Configure API permissions

Add the permissions that match your use case. Choose the permissions from the tables in the [the section called "Permissions"](#) section. Base your selection on your permission scope (all sites or Sites.Selected) and whether you enable ACL crawling.

Content only – Microsoft Graph

- Sites.Read.All

Content only – SharePoint

- `Sites.Read.All`

ACL crawling – Microsoft Graph (additional)

- `Sites.Read.All`
- `User.Read.All`
- `GroupMember.Read.All`

ACL crawling – SharePoint

- `Sites.FullControl.All`

Note

`Sites.FullControl.All` is required for ACL crawling because the SharePoint REST API requires full control permissions to read site-level and item-level permission assignments. If you are using `Sites.Selected`, see [the section called “Step 3b: Grant site-level permissions \(Sites.Selected only\)”](#) for the alternative permission set.

1. In the left navigation of your app registration, choose **API permissions**.
2. Choose **Add a permission**.
3. Choose **Microsoft Graph**.
4. Choose **Application permissions**.
5. Search for and select the required Microsoft Graph permissions for your use case, then choose **Add permissions**.
6. Choose **Add a permission** again.
7. Choose **SharePoint** (under Microsoft APIs).
8. Choose **Application permissions**.
9. Search for and select the required SharePoint permissions for your use case, then choose **Add permissions**.

 Important

Select the **Application permissions** tab, not **Delegated permissions**. Admin-managed setup uses the client credentials flow, which requires application permissions.

Grant admin consent

1. On the **API permissions** page, choose **Grant admin consent for [Your Organization]**.
2. Confirm the consent when prompted.

 Important

Admin consent is required for application permissions. Without it, the application cannot access SharePoint data.

Upload the certificate

1. In the left navigation of your app registration, choose **Certificates & secrets**.
2. Choose the **Certificates** tab.
3. Choose **Upload certificate**.
4. Select the `certificate.cer` file you generated in Step 2.
5. Choose **Add**.

 Note

The Entra portal displays the certificate thumbprint in hexadecimal format. This is different from the base64url-encoded thumbprint you calculated in Step 2. Use the base64url value when you configure the knowledge base in Quick.

Step 3b: Grant site-level permissions (Sites.Selected only)

If you chose `Sites.Selected` as your permission scope, you must explicitly grant your Amazon Quick Entra app access to each SharePoint site. This requires a temporary admin app with `Sites.FullControl.All` permission to call the Microsoft Graph API.

Skip this step if you are using the all-sites permission scope (`Sites.Read.All` or `Sites.FullControl.All`).

Get the Site ID for each SharePoint site

You need the Site ID for each SharePoint site you want to grant access to. To get a Site ID:

1. In your browser, navigate to the SharePoint site (for example, `https://yourcompany.sharepoint.com/sites/SiteName`).
2. Append `/_api/site/id` to the URL and press Enter. For example:
`https://yourcompany.sharepoint.com/sites/SiteName/_api/site/id`
3. The page displays an XML response containing the Site ID (a GUID). Record this value.

Repeat for each site you want to include in the knowledge base.

Create a temporary admin app

The admin app is used only to grant site-level permissions to your Amazon Quick app. You can delete it after completing this step.

1. In the [Microsoft Entra admin center](#), go to **App registrations** and choose **New registration**.
2. For **Name**, enter a descriptive name such as `Quick-SharePoint-PermissionGranter`.
3. For **Supported account types**, select **Accounts in this organizational directory only (Single tenant)**.
4. Leave **Redirect URI** blank and choose **Register**.
5. Record the **Application (client) ID** from the Overview page.
6. Choose **API permissions**, then **Add a permission**.
7. Choose **Microsoft Graph**, then **Application permissions**. Search for and select `Sites.FullControl.All`. Choose **Add permissions**.
8. Choose **Grant admin consent for [Your Organization]** and confirm.

9. Choose **Certificates & secrets**, then **New client secret**. Enter a description, choose an expiration period, and choose **Add**.
10. Record the secret **Value** immediately. This value is only displayed once.

Important

Copy the secret **Value**, not the Secret ID. The Value is the longer string used for authentication.

Get an access token

Use the admin app credentials to retrieve an OAuth token from Microsoft Entra. Replace the *placeholder* values with your admin app's client ID, secret value, and tenant ID.

macOS and Linux (bash)

```
curl -s --location "https://login.microsoftonline.com/TENANT_ID/oauth2/v2.0/token" \  
  --header "Content-Type: application/x-www-form-urlencoded" \  
  --data-urlencode "grant_type=client_credentials" \  
  --data-urlencode "client_id=ADMIN_APP_CLIENT_ID" \  
  --data-urlencode "client_secret=ADMIN_APP_SECRET_VALUE" \  
  --data-urlencode "scope=https://graph.microsoft.com/.default"
```

Windows (PowerShell)

```
$tokenResponse = Invoke-RestMethod `   
  -Uri "https://login.microsoftonline.com/TENANT_ID/oauth2/v2.0/token" `   
  -Method Post `   
  -ContentType "application/x-www-form-urlencoded" `   
  -Body @{   
    grant_type    = "client_credentials"   
    client_id     = "ADMIN_APP_CLIENT_ID"   
    client_secret = "ADMIN_APP_SECRET_VALUE"   
    scope        = "https://graph.microsoft.com/.default"   
  }   
$adminToken = $tokenResponse.access_token
```

The response contains an `access_token` field. Record this value for the next step.

Grant site-level permissions

Use the admin token to grant your Amazon Quick Entra app `fullcontrol` access to each SharePoint site. Replace the *placeholder* values with your Site ID, admin token, and the client app ID and display name from Step 3.

macOS and Linux (bash)

```
curl -s --location "https://graph.microsoft.com/v1.0/sites/SITE_ID/permissions" \
  --header "Content-Type: application/json" \
  --header "Authorization: Bearer ADMIN_TOKEN" \
  --data '{
    "roles": ["fullcontrol"],
    "grantedToIdentities": [{
      "application": {
        "id": "CLIENT_APP_ID",
        "displayName": "CLIENT_APP_NAME"
      }
    }]
  }'
```

Windows (PowerShell)

```
$body = @{
    roles = @("fullcontrol")
    grantedToIdentities = @(
        @{
            application = @{
                id = "CLIENT_APP_ID"
                displayName = "CLIENT_APP_NAME"
            }
        }
    )
} | ConvertTo-Json -Depth 10

Invoke-RestMethod `
  -Uri "https://graph.microsoft.com/v1.0/sites/SITE_ID/permissions" `
  -Method Post `
  -Headers @{
    "Content-Type" = "application/json"
    "Authorization" = "Bearer $adminToken"
  } `
```

```
-Body $body
```

A successful response includes "roles": ["fullcontrol"] and the client app ID in the grantedToIdentities field.

Important

Repeat this command for each SharePoint site you want to include in the knowledge base. Any new sites added in the future also require a separate permission grant.

Clean up

After you have granted permissions to all required sites, you can delete the temporary admin app from the Microsoft Entra admin center. The site-level permissions you granted remain in effect independently of the admin app.

Note

The temporary admin app is used only in your local environment to call the Microsoft Graph API. Amazon Quick never sees or has access to the admin app or its credentials. Only the client app credentials are provided to Amazon Quick when you create the knowledge base.

Step 4: Grant Amazon Quick permission to the KMS key

Amazon Quick needs permission to use the KMS key for signing OAuth assertions. You grant this permission from the Amazon Quick admin console.

Note

This step requires Amazon Quick administrator access (Admin Pro role). If you are not an administrator, ask your Amazon Quick admin to complete this step using the KMS key ARN from Step 1.

Important

If your organization manages its own Amazon Quick IAM service role, the following console steps might not apply. Instead, ensure the role has `kms:Sign` permission on the KMS key ARN from Step 1.

1. In Amazon Quick, choose **Manage account** from the left navigation pane.
2. Under **Permissions**, choose **AWS resources**.
3. On the AWS resources page, scroll to **AWS Key Management Service** and select the checkbox.
4. Choose **Select keys**.
5. In the **Select KMS keys** dialog, enter the KMS key ARN you recorded in Step 1 and choose **Add**.
6. The key ARN appears in the list. Choose **Finish**.
7. Choose **Save** at the bottom of the AWS resources page.

Next steps

After you complete the setup, create the SharePoint Online knowledge base connection in Amazon Quick. For instructions, see [the section called "Create the knowledge base"](#).

Create the knowledge base in Amazon Quick

After you complete the service credentials setup, create the knowledge base in Amazon Quick and provide the credentials. Any enterprise user can complete this step. Amazon Quick administrator access is not required.

You need the following values from [the section called "Set up service credentials"](#). If an administrator completed the setup on your behalf, collect these values before you proceed:

- KMS key ARN (from Step 1)
- Certificate thumbprint in base64url format (from Step 2)
- Entra application (client) ID (from Step 3)
- Directory (tenant) ID (from Step 3)
- SharePoint domain URL (from your Microsoft 365 tenant)

Navigate to Knowledge

1. In Amazon Quick, choose **Knowledge** from the left navigation pane.
2. Under **Set up new knowledge base**, locate **Microsoft SharePoint Online** and choose the **Add** icon.

Select service credentials authentication

The **Create SharePoint knowledge base** wizard opens on the **Authentication method** step. Choose **Have admin access? Connect with service credentials to enable document-level access control**.

If you have previously created a knowledge base using service credentials, the **Connected account** dropdown appears with your most recent connection pre-selected. When you select an existing connection, the SharePoint domain, Tenant ID, and Client ID are displayed below the dropdown for verification. You can use an existing connection or choose **+ Add account** to enter new service credentials.

Enter service credentials

After choosing **+ Add account**, fill in the following fields:

Service credentials form

Field	Description	Example
Connection name	A descriptive name for this SharePoint connection. You see this name when selecting a connection for future knowledge bases.	SharePoint Production
SharePoint domain	Your organization's SharePoint Online domain.	https://contoso.sharepoint.com
Tenant ID	The Microsoft Entra (Azure AD) tenant identifier.	34567890-1234-1234-1234-123456789012
Client ID	The Application (client) ID from your Entra app registration.	12345678-1234-4321-4321-210987654321

Field	Description	Example
KMS key ARN	The ARN of the AWS KMS asymmetric signing key.	arn:aws:kms:us-west-2:123456789012:key/my-key-id
Certificate Thumbprint	The base64url-encoded SHA-1 thumbprint of the certificate uploaded to Entra.	abc123DEF456ghi789EXAMPLE

Fill in each field and choose **Next**.

 Tip

Give the connection a recognizable name. This name appears in the **Connected account** dropdown when you create future knowledge bases.

Choose content

Enter a **Name** and optional **Description** for your knowledge base.

SharePoint site URLs

Enter the URLs of SharePoint sites or subsites you want to sync, such as `https://contoso.sharepoint.com/sites/finance`. Enter one or more site URLs separated by commas and choose **Add Site**. Do not add file or folder URLs here.

You can add up to 350 entries total across both **SharePoint site URLs** and **Specific SharePoint Paths** combined.

Specific SharePoint Paths (optional)

Enter the paths of specific folders, files, lists, libraries, or pages within a SharePoint site. Enter one or more paths separated by commas and choose **Add Path**. The 350-entry limit is shared with **SharePoint site URLs**.

If a site is already added under **SharePoint site URLs**, all content from that site is synced. Adding individual items from the same site does not limit what is synced.

 Important

Do not use browser URLs for paths. To get the correct path, navigate to the item in SharePoint, choose the three dot menu (...), select **Details**, then scroll to **Path** and choose **Copy**.

Example paths:

- Folder: `https://contoso.sharepoint.com/sites/finance/Shared%20Documents/Reports/Q1`
- File: `https://contoso.sharepoint.com/sites/finance/Shared%20Documents/Reports/Q1/data.docx`

After adding your content, choose **Next**.

 Tip

Each site URL and path must be accessible by the Entra app registration. If you are using `Sites.Selected`, ensure per-site permissions have been granted. For more information, see [the section called “Step 3b: Grant site-level permissions \(Sites.Selected only\)”](#).

Configure additional settings

In the final step of the wizard, you can configure optional settings for the knowledge base.

ACL management

Access Control List (ACL) management enables document-level access control for your SharePoint knowledge base. When enabled, Amazon Quick syncs the permission metadata from SharePoint during each crawl. At query time, Quick uses this metadata with a real-time access check. This check ensures that users only see answers from documents they can access in SharePoint.

ACL management is unchecked (disabled) by default. To enable, select the **Control document access with ACLs** checkbox.

When you enable ACL management, the console displays a note that one-time admin consent might be required for real-time access verification. If you are a Microsoft 365 administrator,

choose the **Grant admin consent** link to grant consent directly. If you are not an administrator, ask your Microsoft 365 admin to grant consent. For more information about admin consent, see [the section called "Admin consent"](#).

 Important

ACL management cannot be changed after the knowledge base is created. This setting is immutable. If you need to change this setting, you must create a new knowledge base. Plan your ACL requirements carefully before creating the knowledge base.

 Note

If you enable ACL management, ensure your Entra app has the ACL permission set (User.Read.All, GroupMember.Read.All on Graph, and Sites.FullControl.All on SharePoint). Without these, the knowledge base creation might succeed but ACL enforcement fails during sync.

Filter content by date

Limits which documents are crawled based on their last modified date in SharePoint. The start date defaults to one year before today. You can change or clear the start date, and optionally set an end date.

Multi-media content, file size, and file patterns

- **Visual content in documents** – Extracts and indexes visual elements from supported document formats. Enabled by default.
- **Audio files** – Transcribes and indexes audio files.
- **Video files** – Transcribes and indexes video files.

 Note

Enabling audio and video indexing increases sync time and storage usage. Enable these options only if your SharePoint content includes relevant media files.

Choose **Create** to create the knowledge base.

Initial sync

After you choose **Create**, you are returned to the knowledge base list page. The knowledge base might take a few minutes to finish provisioning. Once creation is complete, an initial sync is automatically triggered. You do not need to start it manually.

Document-level access controls

Admin-managed SharePoint knowledge bases optionally support document-level access control. When enabled, Amazon Quick syncs access control lists (ACLs) from SharePoint during each crawl. The system verifies each user's permissions at query time, so users only see answers from documents they are authorized to access in SharePoint.

How it works

When a user queries an Amazon Quick agent that uses an admin-managed SharePoint knowledge base with ACL management enabled, the system enforces access controls in two stages:

1. **Pre-retrieval filtering** – Amazon Quick performs a semantic search against the vector index to find the most relevant document passages. The system applies access control lists that were synced from SharePoint during the last crawl. This produces a preliminary set of candidate documents.
2. **Real-time verification** – The system verifies the candidate documents in real time by checking the querying user's current access in SharePoint. Only documents the user is currently authorized to access are included in the response.

This two-stage approach provides document-level access control that stays current even when SharePoint permissions change between syncs.

Enable ACL management

ACL management is configured during knowledge base creation in the **Additional settings** step. Select the **Control document access with ACLs** checkbox to enable it.

Important

ACL management cannot be changed after the knowledge base is created. If you need to change this setting, you must create a new knowledge base.

To enable ACL management, your Entra app registration must have the following permissions:

- `User.Read.All` and `GroupMember.Read.All` on Microsoft Graph.
- `Sites.FullControl.All` on the SharePoint resource, or `Sites.Selected` with per-site permissions granted.

For more information about ACL best practices, see [the section called “Best practices for managing ACLs in knowledge bases”](#).

Real-time access verification

The real-time verification stage uses a delegated OAuth flow managed automatically by Amazon Quick. Quick creates and manages a separate Microsoft Entra application specifically for this purpose. No customer configuration is required for this app. It is distinct from both the admin-managed app registration you created during setup and any user-managed OAuth app.

1. A user asks a question in the Quick chat assistant.
2. If the answer involves SharePoint content from an ACL-enabled knowledge base, Quick prompts the user to **Sign in to SharePoint**.
3. The user signs in and accepts the Microsoft consent dialog (if admin consent has not been granted).
4. Quick uses the user's delegated token to verify access to each candidate document in real time.
5. Only documents the user currently has access to in SharePoint are included in the response.

The sign-in is a one-time step. The delegated credentials use a refresh token and last approximately 90 days.

Delegated permissions

The real-time ACL app requests the following delegated permissions:

Real-time ACL – delegated permissions

Permission	Scope	Purpose
Read items in all site collections	<code>Sites.Read.All</code>	Verify user access to SharePoint site content.

Permission	Scope	Purpose
Read your files	Files.Read.All	Verify user access to specific files.
View your basic profile	User.Read	Identify the signed-in user.
Maintain access to data you have given it access to	offline_access	Refresh tokens so users don't need to re-authenticate frequently.

Admin consent

The real-time ACL check uses a separate Microsoft Entra application from the one used in user-managed setup or the admin-managed app registration. If your organization requires admin consent, an administrator must grant consent for each application independently.

When you enable ACL management during knowledge base creation, the Amazon Quick console provides a direct link to grant admin consent. This link is for the real-time ACL application. If you are a Microsoft 365 administrator, you can grant consent directly from the console. Otherwise, share the link with your administrator.

If admin consent is not granted, each user sees the consent dialog on their first query that involves SharePoint content. After accepting, they are not prompted again for approximately 90 days.

For detailed instructions on granting admin consent through the consent dialog or the Microsoft Entra admin center, see [the section called “Grant organization-wide admin consent”](#).

Next steps

For more information about ACL best practices, see [the section called “Best practices for managing ACLs in knowledge bases”](#). For information about creating admin-managed SharePoint knowledge bases, see [the section called “Admin-managed setup”](#).

Manage and troubleshoot admin-managed connections

To edit, share, or delete your integration, see [the section called “Managing existing integrations”](#).

- **Unable to access KMS key** – Verify the KMS key ARN and Region. Confirm the KMS key has been added in the Amazon Quick admin console under **Manage account, AWS resources, AWS Key Management Service**. Confirm the key is enabled and has not been scheduled for deletion. Multi-Region keys are not currently supported.

- **Certificate validation failed** – Verify the thumbprint using the base64url-encoded SHA-1 value from the certificate generation step. Ensure the certificate uploaded to Entra has not expired.
- **ACL not enforced** – Confirm the Entra app has `User.Read.All` and `GroupMember.Read.All` on Microsoft Graph. For the SharePoint resource, confirm the app has `Sites.FullControl.All`. If using `Sites.Selected`, confirm that per-site permissions have been granted for each site in the knowledge base. Re-run a full sync after fixing permissions. For more information about verifying document access, see [the section called “Check document access \(ACL verification\)”](#).
- **Zero items crawled** – The sync completed but no documents were indexed. This typically indicates a permissions issue. Verify the Entra app has the correct API permissions for your permission scope. If using `Sites.Selected`, confirm that per-site permissions have been granted for each site included in the knowledge base (see [the section called “Step 3b: Grant site-level permissions \(Sites.Selected only\)”](#)). Also verify that the SharePoint sites contain content and are accessible.
- **New site not crawled (Sites.Selected)** – If you added a new site URL to the knowledge base but no content from that site is indexed, the Microsoft Graph API permission grant might be missing. Verify that you ran the grant for the new site. Each site requires a separate grant when using `Sites.Selected`. For more information, see [the section called “Grant site-level permissions”](#).
- **Specific paths returning no results** – Verify that you used the SharePoint path, not the browser URL. To get the correct path, navigate to the item in SharePoint, choose the **More options** menu (:), select **Details**, then scroll to **Path** and choose **Copy**. Also verify the path still exists and has not been renamed or moved in SharePoint.
- **Syncs failing after certificate expiry** – If syncs fail across multiple knowledge bases that share the same data source connection, the certificate uploaded to Entra might have expired. Generate a new certificate (see [the section called “Step 2: Generate a self-signed certificate”](#)), upload it to the Entra app registration, and update the connection details. A Amazon Quick administrator can reassign data source ownership through **Manage assets** if the original creator is unavailable. For more information, see [the section called “Sharing data source connections”](#).
- **No results from ACL-enabled knowledge base** – If users receive no results from a knowledge base with ACL management enabled, admin consent for the real-time ACL application might not have been granted. Your tenant might also block user consent. Grant admin consent using the link provided in the Amazon Quick console when ACL management is enabled, or see [the section called “Admin consent”](#).
- **Documents skipped with "File has no ACL" error** – If sync reports show items with status SKIPPED and error type VALIDATION_ERROR with the message "File has no ACL while crawlACL

is true, skipping ingestion," the Entra app registration is missing the required ACL permissions. Verify the app has the correct permissions for your setup. For the required permissions, see [the section called "Permissions"](#).

For additional troubleshooting, including sync monitoring, reports, and ACL verification, see [the section called "Troubleshooting"](#).

Troubleshooting SharePoint knowledge bases

Use this section to diagnose and resolve common issues with SharePoint knowledge base integrations.

Sign-in errors with user-managed setup

User-managed setup sign-in errors

Issue	Cause	Resolution
"Need admin approval" error during sign-in	Your tenant restricts user consent for third-party apps.	Ask a Global Administrator to grant admin consent. For more information, see the section called "Admin consent for Microsoft 365" .
Sign-in window closes without completing	Browser popup blocker or cookie settings.	Allow popups from the Amazon Quick domain and ensure third-party cookies are enabled.
"AADSTS65004: User declined to consent"	User chose Cancel or Deny on the consent dialog.	Retry the sign-in and choose Accept on the consent dialog.

Token refresh with user-managed setup

User-managed setup uses the `offline_access` scope to obtain refresh tokens. Delegated credentials last approximately 90 days. If the refresh token expires or is revoked (for example, by a password change, session revocation, or admin action), syncs fail and you need to re-authenticate.

To re-authenticate:

1. Navigate to the **Knowledge bases** list page in Amazon Quick.

2. Locate the knowledge base that requires a refresh. In the **Actions** column, choose the three-dot menu and choose **Edit integration**.
3. The **Edit integration** dialog opens. You can optionally update the integration name.
4. Choose **Save and Reauthenticate**.
5. A Microsoft sign-in window opens. Complete the sign-in flow using the same credentials that were originally used to create the integration.

Important

When you re-authenticate, sign in using the same Microsoft 365 account that was used during the original setup. Using different credentials might cause issues with connected resources. The new account might have different SharePoint permissions than the original account.

Common errors

Common SharePoint knowledge base errors

Issue	Cause	Resolution
Unable to access KMS key	The KMS key ARN is incorrect, the key does not exist in the specified Region, or the KMS key has not been added in the Amazon Quick admin console.	Verify the KMS key ARN and Region. Confirm the KMS key has been added under Manage account, AWS resources , AWS Key Management Service . Confirm the key is enabled and has not been scheduled for deletion. Multi-Region keys are not supported.
Invalid Tenant ID	The Tenant ID is malformed or does not match a valid Entra tenant.	Copy the Tenant ID directly from the Entra admin center: Microsoft Entra ID, Overview .
Invalid Client ID	The Client ID is incorrect or the app registration does not exist.	Copy the Application (Client) ID from App Registrations , your app, Overview .

Issue	Cause	Resolution
Certificate validation failed	The thumbprint does not match the certificate in Entra, or the certificate has expired.	Verify the thumbprint using the base64url-encoded SHA-1 value from the certificate generation step. Ensure the certificate has not expired.
Unable to connect to SharePoint	One or more credentials are incorrect, or the Entra app is missing required API permissions or admin consent.	Review all field values. Confirm the Entra app has the correct API permissions and that admin consent has been granted.
Error loading the file picker with user-managed setup	The file picker is unable to establish a connection to SharePoint on behalf of the signed-in user.	Close the error message and choose Add content again to retry. If the error persists, close the creation dialog, reopen the flow, and sign in to SharePoint again.

Check document access (ACL verification)

If you enabled ACL management, you can verify document-level access control for individual items using the Permission Checker in sync reports. For detailed instructions, see [the section called "Check document access \(ACL verification\)"](#).

If you see "No access control list found" for a SharePoint document:

- Verify **Enable ACL management** is selected in the knowledge base's additional settings.
- Confirm the Entra app has `User.Read.All` and `GroupMember.Read.All` on Microsoft Graph.
- Confirm the Entra app has `Sites.FullControl.All` on the SharePoint resource (or `Sites.Selected` with per-site permissions granted).
- Re-run a full sync after fixing permissions.

Sync monitoring and reports

For detailed information about sync schedules, sync activity, sync reports, filtering items, and downloading reports, see [the section called "Sync reports and observability"](#).

Tip

To trigger a sync on demand, choose **Sync now** on the knowledge base detail page.

Next steps

After you resolve the issue, verify your fix by running a sync. For more information about sync monitoring, see [the section called “Sync reports and observability”](#).

Microsoft Teams integration

Use the Microsoft Teams connector to send messages, manage channels, schedule meetings, and manage team collaboration directly in Amazon Quick through natural language.

Amazon Quick supports multiple authentication methods for Microsoft Teams. Choose the method that best fits your organization's security requirements.

- **Default OAuth app** – Uses an AWS-managed OAuth application. No additional credentials are needed. Users authenticate directly with their Microsoft account.
- **Custom OAuth app** – Uses a customer-managed application registered in Microsoft Entra. This option gives your organization full control over the OAuth configuration. Users authenticate on behalf of a signed-in user (delegated permissions).
- **Service-to-Service OAuth** – Uses client credentials for server-to-server authentication without user interaction (application permissions). Suitable for automated workflows.

For more information about the authentication methods that Amazon Quick supports, see [the section called “Authentication methods”](#).

Before you begin

Make sure you have the following before you set up the integration.

- A Microsoft 365 account with Teams access.
- For **Custom OAuth app** or **Service-to-Service OAuth**: Access to the [Microsoft Entra admin center](#) on the Microsoft website with at least Application Developer permissions.
- For subscription requirements, see [the section called “Set up integrations in the console”](#).

Configure Microsoft Entra

If you are using **Default OAuth app** authentication, skip this section and proceed to [the section called "Setting up the connector in Amazon Quick"](#).

Before you configure Amazon Quick, create an app registration in Microsoft Entra. Complete all of the following steps in Entra before moving to the Amazon Quick console.

For more information about app registrations, see [Register an application with the Microsoft identity platform](#) in the Microsoft documentation.

Register the application

1. Open the [Microsoft Entra admin center](#).
2. In the left navigation, choose **Entra ID**, then choose **App registrations**.
3. Choose **New registration**.
4. For **Name**, enter a descriptive name for your integration.
5. For **Supported account types**, choose **Accounts in this organizational directory only**.
6. For **Redirect URI**, select **Web** and enter `https://{region}.quicksight.aws.amazon.com/sn/oauthcallback`. Replace `{region}` with the AWS Region where your Amazon Quick instance is deployed.
7. Choose **Register**.
8. On the overview page, copy the **Application (client) ID** and **Directory (tenant) ID**. You need these values for the Amazon Quick configuration.

Create a client secret

Amazon Quick needs a client secret to authenticate with Microsoft Entra. This secret acts as a password for the app registration.

1. From your app registration, choose **Certificates & secrets**.
2. Choose **New client secret**.
3. Enter a description and choose an expiration period.
4. Choose **Add**.
5. Copy the **Value** immediately. This value is only displayed once.

⚠ Important

Copy the secret **Value**, not the Secret ID. The Value is the longer string used for authentication.

Configure API permissions

Microsoft Graph supports two permission types for this integration. Delegated permissions allow the app to act on behalf of a signed-in user. Application permissions allow the app to act without a signed-in user. For more information, see [Overview of Microsoft Graph permissions](#) in the Microsoft documentation.

1. From your app registration, choose **API permissions**.
2. Choose **Add a permission**, then choose **Microsoft Graph**.
3. Choose **Delegated permissions** or **Application permissions** based on your authentication method, and add the permissions from the appropriate table below.
4. Choose **Grant admin consent for [your tenant name]** to approve the permissions.

For user authentication (delegated permissions):

Add the following as Delegated permissions in your Entra app registration. For the full permissions reference, see [Microsoft Graph permissions reference](#) in the Microsoft documentation.

Teams action integration – delegated permissions

Permission	Description
Calendars.ReadWrite	Allows the app to read and write events in user calendars on behalf of the signed-in user.
Channel.Create	Allows the app to create channels in any team on behalf of the signed-in user.
Channel.ReadBasic.All	Allows the app to read channel names and descriptions on behalf of the signed-in user.
ChannelMessage.Read.All	Allows the app to read all channel messages on behalf of the signed-in user.

Permission	Description
<code>ChannelMessage.ReadWrite</code>	Allows the app to read, send, update, and delete messages in channels on behalf of the signed-in user.
<code>ChannelMessage.Send</code>	Allows the app to send messages in channels on behalf of the signed-in user.
<code>ChannelSettings.ReadWrite.All</code>	Allows the app to read and write the settings of all channels on behalf of the signed-in user.
<code>Chat.Create</code>	Allows the app to create chats on behalf of the signed-in user.
<code>Chat.ReadWrite</code>	Allows the app to read and write the signed-in user's chat messages.
<code>ChatMessage.Send</code>	Allows the app to send chat messages on behalf of the signed-in user.
<code>Files.Read</code>	Allows the app to read the signed-in user's files.
<code>offline_access</code>	Allows the app to refresh access tokens without requiring the user to sign in again. This reduces how often users need to re-authenticate.
<code>OnlineMeetingRecording.Read.All</code>	Allows the app to read all recordings of online meetings on behalf of the signed-in user.
<code>OnlineMeetings.ReadWrite</code>	Allows the app to read and create online meetings on behalf of the signed-in user.
<code>OnlineMeetingTranscript.Read.All</code>	Allows the app to read all transcripts of online meetings on behalf of the signed-in user.
<code>People.Read</code>	Allows the app to read a list of people relevant to the signed-in user.
<code>Presence.ReadWrite</code>	Allows the app to read and write the presence information of the signed-in user.

Permission	Description
<code>Sites.Read.All</code>	Allows the app to read documents and list items in all site collections on behalf of the signed-in user.
<code>Team.Create</code>	Allows the app to create teams on behalf of the signed-in user.
<code>Team.ReadBasic.All</code>	Allows the app to read the names and descriptions of teams on behalf of the signed-in user.
<code>TeamMember.ReadWrite.All</code>	Allows the app to add and remove members from all teams on behalf of the signed-in user.
<code>User.Read</code>	Allows the app to read the signed-in user's profile.
<code>User.Read.All</code>	Allows the app to read the full set of profile properties of all users on behalf of the signed-in user.
<code>User.ReadBasic.All</code>	Allows the app to read a basic set of profile properties of all users on behalf of the signed-in user.

For service authentication (application permissions):

Add the following as Application permissions in your Entra app registration.

Teams action integration – application permissions

Permission	Description
<code>Chat.Read.All</code>	Allows the app to read all chat messages in your organization without a signed-in user.
<code>Chat.Send</code>	Allows the app to send chat messages without a signed-in user.
<code>Team.ReadBasic.All</code>	Allows the app to read the names and descriptions of all teams without a signed-in user.
<code>Channel.ReadBasic.All</code>	Allows the app to read all channel names and descriptions without a signed-in user.

Permission	Description
<code>ChannelMessage.Read.All</code>	Allows the app to read all channel messages without a signed-in user.
<code>ChannelMessage.Send</code>	Allows the app to send messages to any channel without a signed-in user.
<code>ChannelMember.ReadWrite.All</code>	Allows the app to add and remove members from all channels without a signed-in user.
<code>TeamMember.ReadWrite.All</code>	Allows the app to add and remove members from all teams without a signed-in user.
<code>User.Read.All</code>	Allows the app to read the full set of profile properties of all users without a signed-in user.
<code>OnlineMeetingTranscript.Read.All</code>	Allows the app to read all transcripts of online meetings without a signed-in user.

Important

With service authentication, all actions execute as the service account. Any user with access to this integration can perform actions across all teams and channels that the service account can access. Scope the application permissions appropriately for your organization's security requirements.

Record your credentials

Before leaving the Microsoft Entra admin center, confirm you have the following values. You need them for the Amazon Quick configuration.

Required credentials from Microsoft Entra

Value	Where to find it
Application (client) ID	App registration overview page

Value	Where to find it
Directory (tenant) ID	App registration overview page
Client secret value	Certificates & secrets page

Setting up the connector in Amazon Quick

Connect from the Available tab

If you want to use Default OAuth app authentication, you can connect directly from the **Available** tab without additional configuration.

1. In the Amazon Quick console, choose **Connectors**.
2. On the **Available** tab, find **MSTeams** and choose **Connect**.
3. Complete the Microsoft sign-in flow and grant the requested permissions.

To configure a connector with Custom OAuth app or Service-to-Service OAuth instead, use the **Create for your team** tab as described below.

Create from the Create for your team tab

After you complete any required Entra configuration, create the connector in Amazon Quick.

1. In the Amazon Quick console, choose **Connectors**.
2. Choose the **Create for your team** tab.
3. Find and choose **Microsoft Teams**.

Note

If a Microsoft Teams connector already exists, a dialog appears with your existing connectors. To use an existing connector, choose it. To create a new one, choose **No, create new**.

4. Enter a **Name** for your connector. Optionally, choose **+ Add Description** to add a description.
5. For **Connection type**, choose **Public network**.

6. For **OAuth Configuration**, choose one of the following authentication methods and configure the required fields.

a. For **Default OAuth app**:

No additional credentials are needed. Choose **Next** to continue.

b. For **Custom OAuth app** (user authentication with delegated permissions), configure the following fields:

- **Base URL** (Optional) – The Microsoft Graph API base URL. Example: `https://graph.microsoft.com/v1.0`
- **Client ID** – The Application (client) ID from your Entra app registration.
- **Client secret** – The client secret value from your Entra app registration.
- **Token URL** – The token endpoint. Example: `https://login.microsoftonline.com/{tenant-id}/oauth2/v2.0/token`
- **Authorization URL** – The authorization endpoint. Example: `https://login.microsoftonline.com/{tenant-id}/oauth2/v2.0/authorize`
- **Redirect URL** – Pre-filled with the Amazon Quick callback URL.

c. For **Service-to-Service OAuth** (service authentication with application permissions), configure the following fields:

- **Base URL** (Optional) – The Microsoft Graph API base URL. Example: `https://graph.microsoft.com/v1.0`
- **Client ID** – The Application (client) ID from your Entra app registration.
- **Client secret** – The client secret value from your Entra app registration.
- **Token URL** – The token endpoint. Example: `https://login.microsoftonline.com/{tenant-id}/oauth2/v2.0/token`

 Note

The scope for the client credentials token request (`https://graph.microsoft.com/.default`) is set automatically by Amazon Quick. You do not need to configure it manually.

7. Choose **Next**.

8. If you chose **Default OAuth app** or **Custom OAuth app**, a Microsoft authorization window opens. Review the requested permissions and choose **Accept**.

If you see an error instead of the consent dialog, your organization might restrict third-party app access. See [the section called "Admin consent for Microsoft 365"](#).

9. On the **Review** page, review the available actions for the connector. Choose **Next**.
10. On the **Publish** page, choose who can access the connector. You can enable access for everyone in your organization or search for specific teams or groups.
11. Choose **Publish**.

Available actions

After you set up the integration, the following actions are available.

Microsoft Teams available actions

Category	Action	Description
Chats	List Chats	View your chat conversations.
Chats	Create Chat	Start a new chat conversation.
Chats	Send Chat Message	Send a message in a chat.
Teams	List Teams	View teams you're a member of.
Teams	List Team Members	View members of a team.
Teams	Add Team Member	Add a member to a team.
Channels	List Channels	View channels in a team.
Channels	Create Channel	Create a new channel in a team.
Channels	List Channel Messages	View messages in a channel.
Channels	Send Channel Message	Post a message to a channel.
Channels	List Channel Members	View members of a channel.

Category	Action	Description
Channels	Add Channel Member	Add a member to a channel.
Users	List Users	View users in your organization.
Meetings	List Online Meetings	View your scheduled online meetings.
Meetings	Create Online Meeting	Schedule a new online meeting.
Meetings	List Meeting Transcripts	View transcripts from meetings.
Calendar	List Calendar Events	View events on your calendar.
Calendar	Create Calendar Event	Create a new calendar event.

Manage and troubleshoot

To edit, share, or delete your integration, see [the section called “Managing existing integrations”](#).

Authentication issues

- **Incorrect app registration** – Verify the app registration in Microsoft Entra includes the required API permissions and that admin consent has been granted.
- **Expired client secret** – Check if the client secret has expired in **Certificates & secrets** and generate a new one if needed.
- **Incorrect redirect URI** – Verify the redirect URI in Microsoft Entra matches `https://{region}.quicksight.aws.amazon.com/sn/oauthcallback`.

Common error messages

- **Access denied. You do not have permission to perform this action** – The authenticated user does not have the required permissions. Contact your administrator to verify and grant appropriate permissions.
- **AADSTS50020: User account from identity provider does not exist in tenant** – The user account is not configured in the correct Microsoft Entra tenant. Verify the user account exists in the tenant that matches the Directory (tenant) ID in your app registration.

Admin consent for Microsoft 365

When you use the **Default OAuth app** authentication method, Amazon Quick uses an AWS-managed application to access Microsoft Teams on behalf of the signed-in user. Most users can complete setup without any extra steps. However, if your Microsoft 365 tenant restricts third-party app access, a Microsoft 365 administrator must grant one-time consent before users can connect.

If you see an error when you sign in during connector setup, your organization might restrict third-party app access. Share the following information with your Microsoft 365 administrator:

- **What to do:** Grant admin consent for the Amazon Quick Microsoft Teams integration application.
- **Why:** Amazon Quick needs delegated access to Teams channels, chats, meetings, and calendar data to perform actions on behalf of users.

An administrator can grant consent in one of the following ways:

- **Through the consent dialog** – A Global Administrator or Privileged Role Administrator initiates the connector setup flow. In the Microsoft sign-in dialog, they select the **Consent on behalf of your organization** check box and choose **Accept**.
- **Through the Microsoft Entra admin center** – Sign in to the [Microsoft Entra admin center](#) on the Microsoft website. Choose **Enterprise applications**, locate the Amazon Quick application, choose **Permissions**, and choose **Grant admin consent for *Your Organization***.

After consent is granted, any user in your organization can connect without being prompted for individual consent.

Note

To check whether your tenant restricts user consent, go to the Microsoft Entra admin center and choose **Enterprise applications, Consent and permissions, User consent settings**. If the setting is **Do not allow user consent**, an administrator must grant consent before users can use the connector.

Model Context Protocol (MCP) integration

Model Context Protocol (MCP) is an open standard that defines how AI applications communicate with external tools and data sources. MCP uses a client-server architecture. AI applications act

as clients that connect to MCP servers. Each MCP server exposes a set of tools. These tools are structured operations that the AI application can invoke to perform tasks, such as querying databases, calling APIs, or interacting with third-party services.

With MCP integration in Amazon Quick, you can connect to remote MCP servers so that your AI assistant can use the tools that those servers provide. For example, you can connect to an MCP server that provides access to your project management system. This connection allows the assistant to create tickets, look up issues, or update statuses as part of a conversation. Because MCP is an open standard, you can connect to any compatible server without building custom integrations for each tool.

You can connect to MCP servers that are reachable over the public internet. You can also connect to private MCP servers that are reachable from a virtual private cloud (VPC) in your AWS account by using a Amazon Quick VPC connection. The MCP server can be in the VPC, in another VPC reachable through peering or transit gateway, or on-premises reachable through AWS Direct Connect or VPN. For information about creating a VPC connection in Amazon Quick, see [Configuring VPC connections in Amazon Quick Sight](#).

Important

When you connect to a private MCP server through a VPC connection, you can configure two independent VPC connections:

- **Resource-server VPC connection** (`VpcConnectionArn`) – Carries traffic between Amazon Quick and the MCP server endpoint (tool invocations, capability discovery).
- **Auth-server VPC connection** (`AuthVpcConnectionArn`) – Carries OAuth traffic between Amazon Quick and the authorization server (token requests, authorization code exchange, Dynamic Client Registration, and metadata discovery).

You can set one, both, or neither connection depending on your network topology. If you do not set an auth-server VPC connection, OAuth endpoints must be reachable over the public internet (the default behavior). If your OAuth provider is reachable only from within a VPC, set an auth-server VPC connection. The resource-server and auth-server VPC connections are independent. They can point to different VPCs. MCP servers that do not require authentication are also supported.

Supported capabilities

MCP integration registers MCP server tools as actions in Amazon Quick.

Connector

Each tool that an MCP server exposes registers as an action that your AI assistant can invoke during conversations. The integration secures these connections by using Proof Key for Code Exchange (PKCE) with the S256 challenge method and Resource Indicators (RFC 8707) to bind access tokens to specific MCP servers.

Before you begin

Before you set up MCP integration, make sure that you have the following:

- An MCP server endpoint with appropriate access.
- Authentication credentials for the MCP server, if required. For more information, see [the section called “Prepare MCP server setup and authentication”](#).
- An Amazon Quick Enterprise subscription.
- If you are connecting to a private MCP server, an active Amazon Quick VPC connection that has network access to the MCP server. For more information, see [Configuring VPC connections in Amazon Quick Sight](#).

Note

MCP integration supports remote servers only. HTTP streaming is preferred over Server-Sent Events (SSE). Local stdio connections are not supported.

Prepare MCP server setup and authentication

When you connect to an MCP server, Amazon Quick uses OAuth 2.0 Protected Resource Metadata (RFC 9728) to automatically discover authorization server information. The client sends an initial unauthenticated request to the MCP server. If the server responds with a 401 status that contains a WWW-Authenticate header with a resource_metadata URL, then Amazon Quick uses that URL to fetch the metadata document. If the header is not present, then Amazon Quick falls back to the well-known URI at the server root.

When an auth-server VPC connection is configured, Amazon Quick performs all OAuth discovery over that connection instead of the public internet. This includes fetching the `/.well-known/oauth-authorization-server` document, the `resource_metadata` URL from the MCP server's `WWW-Authenticate` header, and the `registration_endpoint` for Dynamic Client Registration (DCR). DNS resolution for the authorization server hostname must be reachable from the subnets and DNS resolver endpoints that are associated with the auth-server VPC connection.

If the authorization server supports Dynamic Client Registration (DCR), then Amazon Quick automatically registers itself by using the discovered `registration_endpoint` from the authorization server metadata. No manual credential configuration is required. Amazon Quick supports both confidential and public client flows. DCR applies regardless of the authentication method that you choose.

If the authorization server does not support DCR, then you must manually provide credentials. Choose the authentication method that matches your MCP server requirements.

User authentication (OAuth)

Gather the following information from your MCP server configuration:

- **Client ID** – The OAuth client ID.
- **Public OAuth client** (Optional) – Select this checkbox if your OAuth application uses public client configuration (no client secret).
- **Client Secret** – The OAuth client secret.
- **Token URL** – The OAuth token endpoint.
- **Authorization URL** – The OAuth authorization endpoint.
- **Redirect URL** – The OAuth redirect URI.

Note

Select **Public OAuth client** if your MCP server's OAuth application does not use a client secret. This is common for public clients that rely solely on PKCE for security. When this option is selected, the **Client Secret** field is not required.

Service authentication (Service-to-Service)

Gather the following information from your MCP server configuration:

- **Client ID** – The service client ID.

- **Client Secret** – The service client secret.
- **Token URL** – The service token endpoint.

No authentication

If the MCP server does not require authentication, no credentials are needed. Select this option for MCP servers that allow unauthenticated access.

Set up MCP integration

After you prepare your MCP server configuration and authentication credentials, create your MCP integration.

1. In the Amazon Quick console, choose **Connectors**.
2. Choose the **Create for your team** tab.
3. Find and choose **Model Context Protocol (MCP)**.
4. On the **Create Integration** page, enter the integration details:
 - **Name** – A descriptive name for your MCP integration.
 - **Description** (Optional) – The purpose of the integration.
 - **MCP server endpoint** – The URL of the MCP server.
 - **Connection type** – Choose how Amazon Quick connects to the MCP server (resource server):
 - **Public network** – Use this option for MCP servers that are reachable over the public internet.
 - A named VPC connection – Use this option for private MCP servers. The dropdown lists the VPC connections that are configured on your Amazon Quick account. Choose the connection that has network access to your MCP server. If you don't see your VPC connection in the list, confirm that it is fully provisioned and active. For more information, see [Configuring VPC connections in Amazon Quick Sight](#).
 - **Auth server connection type** – Choose how Amazon Quick connects to the OAuth authorization server. This setting is independent of the resource-server connection type.
 - **Public network** (default) – Use this option when the OAuth endpoints (authorization, token, and registration endpoints) are reachable over the public internet. This preserves the default behavior for existing integrations.
 - A named VPC connection – Use this option when the OAuth provider is reachable only from within a VPC. Choose the VPC connection that has network access to the

authorization server. The auth-server VPC connection can be the same as or different from the resource-server VPC connection. For example, your MCP server and OAuth provider might reside in separate VPCs.

5. Choose **Next**.
6. Select the authentication method (user, service, or no authentication).
7. Enter the appropriate configuration details.
8. Choose **Create and continue**.
9. Review the integration details.
10. Choose **Next**.
11. Share the integration with other users if needed.

After you create your MCP integration, Amazon Quick discovers the available tools and registers them as actions.

Review integration

After you configure authentication, review the MCP integration capabilities:

1. The system connects to the MCP server and discovers available capabilities.
2. Review the list of available actions and tasks that the MCP server provides.
3. Confirm the integration configuration and capabilities.

Capability discovery

During the connection process that is described in [the section called “Prepare MCP server setup and authentication”](#), Amazon Quick also discovers and registers the tools that the MCP server provides. After discovery completes, each tool is listed as an action that you can review and turn on.

Manage MCP integrations

To edit, share, or delete your integration, see [the section called “Managing existing integrations”](#).

MCP Sync

As external MCP servers add new tools, update descriptions, and evolve their capabilities, MCP Sync keeps your connectors current. When an MCP server adds, removes, or modifies tools, MCP Sync updates your connector to reflect those changes without requiring you to reconnect.

Amazon Quick supports two connector types with different sync behaviors:

Built-in MCP connectors

Built-in connectors sync automatically. When the MCP server updates its tools, Amazon Quick detects the change and applies it without any action from you. Auto-sync is mandatory for built-in connectors and cannot be turned off. This prevents action failures caused by schema changes or removed tools on the server side. Built-in connectors sync automatically twice a month, on the first and third Monday of each month at 9:00 AM Pacific time.

Custom MCP connectors

Custom connectors require the connector owner to trigger a sync manually. To sync a custom connector, open the connector details page, and then choose **Sync**.

Note

Triggering a sync initiates a re-authorization flow. If the MCP server now requires OAuth scopes that were not granted during the original connection, the sync fails. Confirm that all required scopes are authorized before you sync.

Security considerations

When you use an auth-server VPC connection with MCP integrations, the following security behaviors apply.

Data in transit

When an auth-server VPC connection (`AuthVpcConnectionArn`) is configured, all OAuth credential exchange, including authorization codes, refresh tokens, and ID tokens, transits your VPC over TLS. Tokens do not traverse the public internet on the auth-server path and remain within AWS internal infrastructure.

Logging

The auth-server VPC connection does not introduce new AWS CloudTrail event types. Existing CloudTrail coverage on `CreateActionConnector`, `UpdateActionConnector`, and `PatchActionConnector` API operations continues to apply. The `AuthVpcConnectionArn` field appears in the request body of those existing events when it is set.

Troubleshooting

Use the following guidance to diagnose and resolve common issues when you create or use an MCP integration in Amazon Quick.

Connector creation issues

- **Connector passes discovery but fails at publish with `Creation failed`** – This error usually means that one or more tool definitions in your MCP server's `tools/list` response contain an invalid `inputSchema`. Amazon Quick validates each tool's `inputSchema` against JSON Schema Draft 7 or later during the publish phase.

The most common cause is the deprecated Draft 3 syntax, where `required` is a boolean inside a property definition (for example, `"required": true`). In JSON Schema Draft 7 and later, `required` must be an array of property names at the schema root, as a sibling of `properties`.

Update your tool definitions to use the correct format and redeploy your MCP server. After you redeploy, delete the failed connector and create a new one to retrigger discovery and publish.

The following example shows the incorrect Draft 3 syntax:

```
{
  "type": "object",
  "properties": {
    "logNumber": {
      "type": "string",
      "description": "The permit log number",
      "required": true
    }
  }
}
```

The following example shows the correct Draft 7 syntax:


```
{
  "type": "object",
  "properties": {
    "logNumber": {
      "type": "string",
      "description": "The permit log number"
    }
  },
  "required": ["logNumber"]
}
```

Some MCP framework libraries and code generators emit Draft 3 syntax by default. Check your framework documentation for the option that selects the JSON Schema output version. You can also validate your schemas with any JSON Schema Draft 7 validator before you deploy. For more information about `inputSchema` requirements, see [Tools](#) in the Model Context Protocol specification.

You might also notice that connector creation hangs for two to five minutes before it fails. The elapsed time reflects internal retries during the publish phase, not a network timeout. The fix is the same: check your tool `inputSchema` definitions for JSON Schema Draft 7 compliance.

VPC connection issues

- **VPC connection is not listed in the Connection type dropdown, or MCP creation fails immediately after you select it** – The VPC connection might appear in the dropdown before provisioning is complete. Wait until the VPC connection shows as available on the Amazon Quick admin page before you create an MCP integration against it. For more information, see [Configuring VPC connections in Amazon Quick Sight](#).
- **MCP creation fails because the MCP server hostname cannot be resolved** – Amazon Quick does not use the default VPC DNS resolver for MCP integrations. You must populate the **DNS resolver endpoints** field on the VPC connection with Route 53 Resolver inbound endpoint IP addresses that can resolve your MCP server hostname. This requirement applies to both private hostnames and to public hostnames that you want to resolve to a private address from within the VPC, such as endpoints fronted by AWS PrivateLink or a Route 53 private hosted zone. Without these resolver endpoints, the MCP server hostname cannot be resolved and integration creation fails. For more information, see [Configuring VPC connections in Amazon Quick Sight](#).

- **MCP server is unreachable from the VPC connection** – Traffic from Amazon Quick to your MCP server originates from the subnets that you selected during VPC connection setup. Confirm that route tables, network ACLs, and security groups allow traffic between those subnets and the MCP server endpoint. For general VPC networking guidance, see [Configuring VPC connections in Amazon Quick Sight](#).
- **Authorization server hostname cannot be resolved from the auth-server VPC connection** – The same DNS resolver requirement that applies to the resource-server VPC connection also applies to the auth-server VPC connection. You must populate the **DNS resolver endpoints** field on the auth-server VPC connection with Route 53 Resolver inbound endpoint IP addresses that can resolve the authorization server hostname. If the resource server and authorization server are in different VPCs, each VPC connection needs its own DNS resolver endpoints. For more information, see [Configuring VPC connections in Amazon Quick Sight](#).
- **Authorization server is unreachable from the auth-server VPC connection** – Traffic from Amazon Quick to the authorization server originates from the subnets that are associated with the auth-server VPC connection. Confirm that route tables, network ACLs, and security groups in those subnets allow outbound traffic to the authorization server endpoints. The resource-server and auth-server VPC connections might use different subnets in different VPCs, so verify networking requirements independently for each connection.
- **Authentication fails when you use a private MCP server** – If the OAuth endpoints that the MCP server advertises are reachable only from within a VPC, you must configure an auth-server VPC connection. Without an auth-server VPC connection, Amazon Quick attempts to reach the OAuth endpoints over the public internet, and requests fail if those endpoints are not publicly accessible.

To resolve this issue, edit the MCP integration and set the **Auth server connection type** to the VPC connection that has network access to the authorization server. After you save the change, Amazon Quick routes all OAuth traffic (token requests, authorization code exchange, metadata discovery, and Dynamic Client Registration) through that VPC connection.

If an auth-server VPC connection is already configured and authentication still fails, confirm the following:

- The **DNS resolver endpoints** on the auth-server VPC connection can resolve the authorization server hostname.
- Route tables, network ACLs, and security groups in the auth-server VPC connection's subnets allow outbound traffic to the authorization server endpoints.

- The OAuth URLs returned in the MCP server's protected resource metadata match the endpoints that the auth-server VPC connection can reach.

Microsoft Entra ID issues

If your MCP server uses Microsoft Entra ID as the authorization server, the following errors are common during connector setup. These issues apply to both public network connections and private MCP servers that you reach through a VPC connection.

- **AADSTS9010010 – the v2.0 endpoint rejects the resource parameter** – Amazon Quick sends a resource parameter on OAuth requests as required by the MCP specification (RFC 8707). The Entra ID v2.0 endpoint rejects requests that include both a resource parameter and scope values. To resolve this error, configure your app registration to use the Entra ID v1.0 OAuth endpoints and set `accessTokenAcceptedVersion` to 2 in the app manifest.
- **AADSTS90009 – application requests a token for itself** – When you use user authentication (authorization code flow) and the OAuth client and the MCP resource resolve to the same Entra ID application, Entra ID blocks the request. Create two separate app registrations: one client app for Amazon Quick and one resource app for the MCP server. This issue does not affect service authentication (client credentials flow), where a single app registration works.

Snowflake MCP server issues

If your MCP server is hosted on Snowflake, the following issue is common during connector setup. This issue applies to both public network and VPC connection configurations.

- **Connection fails or returns unexpected errors when the endpoint URL does not include the correct Region** – Snowflake MCP server endpoint URLs are Region-specific. If the endpoint URL omits the AWS Region or defaults to a different Region than where your Snowflake account is deployed, the connection fails.

Verify that the **MCP server endpoint** URL that you provide during integration setup includes the Region where your Snowflake account is deployed. The general endpoint format is:

```
https://account-identifier.region.snowflakecomputing.com/api/v2/cortex/mcp
```

If you connect through AWS PrivateLink, use the following format:

```
https://account-identifier.region.privatelink.snowflakecomputing.com/api/v2/cortex/mcp
```

Replace *account-identifier* with your Snowflake account identifier and *region* with the AWS Region where your Snowflake account is deployed (for example, *us-east-1*). For the correct endpoint URL format for your account, see your Snowflake documentation.

If you encounter an issue that is not covered in the preceding sections, contact AWS Support through the Amazon Quick console or AWS Support Center.

Limitations

When you use MCP integrations in Amazon Quick, be aware of the following limitations:

- MCP operations have a fixed 60-second timeout. Operations that exceed this limit automatically fail with an HTTP 424 error.
- For MCP servers that you reach through a VPC connection, DNS resolution for the authorization server hostname must be reachable from the auth-server VPC connection's subnets and DNS resolver endpoints. This is the same constraint that applies to the resource-server VPC connection for the MCP server hostname. If you do not configure an auth-server VPC connection, the OAuth endpoints must be reachable over the public internet.
- Custom HTTP headers are not supported in MCP operations. Only standard system headers are transmitted.
- For custom MCP connectors, tool lists do not update automatically. To pick up server-side tool changes, open the connector details page and choose **Sync**. Built-in MCP connectors sync automatically. For more information, see [MCP Sync](#).
- Amazon Quick supports a maximum of 100 tools per MCP server connection. If the MCP server exposes more than 100 tools, only the first 100 are registered.
- Connector creation might fail if the Amazon Quick callback URI is not allow-listed by third-party providers.
- Server connectivity issues result in immediate failure without retry attempts.
- Step-up authorization is not supported. If an MCP server requires additional scopes after the initial authorization (HTTP 403 with `insufficient_scope`), then you must re-authorize the entire connection. Incremental permission upgrades are not available.
- Scope handling has the following limitations:

- Amazon Quick does not extract the scope parameter from the server's initial 401 WWW-Authenticate challenge. Amazon Quick determines scopes from the Protected Resource Metadata document instead.
- When the metadata does not specify supported scopes, Amazon Quick applies default scopes rather than omitting them. This behavior might cause authentication failures with servers that do not recognize the default scopes.
- Only Dynamic Client Registration (DCR) is supported for automatic client registration. Client ID Metadata Documents are not supported.
- Well-known URI discovery uses the server root path only. Path-specific metadata locations (path-insertion discovery) are not supported. This limitation might prevent discovery of servers that serve metadata only at path-specific URIs.

Monday.com integration

With Monday.com integration in Amazon Quick, you can manage project boards, track items, and coordinate team workflows through MCP server connectivity. This integration provides action capabilities for project management and team collaboration operations.

What you can do

Monday.com integration provides connector capabilities through MCP server connectivity:

- Create and manage project boards
- Add and update board items and tasks
- Manage team assignments and workloads
- Track project timelines and deadlines
- Update item status and progress
- Generate reports and analytics

Available tools

The Monday.com MCP server typically provides these tools:

- `create_board` - Create new project boards
- `get_board` - Get board information

- `create_item` - Create new board items
- `update_item` - Update item details
- `list_items` - List board items
- `create_column` - Create board columns
- `update_column_value` - Update column values
- `search_boards` - Search for boards

Note

The specific tools and capabilities available through this MCP server may change over time. For the most current information about supported tools, features, and implementation details, check the official Monday.com documentation and MCP server repository.

Setting up Monday.com integration

Monday.com integration uses MCP server connectivity to provide action capabilities. For detailed setup instructions, see [the section called “MCP integration”](#).

You'll need:

- Monday.com account with appropriate board permissions
- Monday.com API token for authentication

Compatibility

Monday.com integration supports:

- **Chat Agents:** Yes
- **Flows:** Yes
- **Knowledge Base:** No

Moodys GenAI Ready Data integration

With the Moodys GenAI Ready Data connector, you can access credit ratings, financial research, and risk analytics data directly in Amazon Quick through natural language.

Amazon Quick supports two authentication methods for Moodys GenAI Ready Data. Choose the method that best fits your organization's security requirements.

- **Default OAuth app** – Uses an AWS-managed OAuth application. No additional credentials are needed. Users authenticate directly with their Moody's account.
- **Custom OAuth app** – Uses a customer-managed OAuth application. This option gives your organization full control over the OAuth configuration.

For more information about the authentication methods that Amazon Quick supports, see [the section called "Authentication methods"](#).

Before you begin

Make sure that you have the following before you set up the integration.

- An active Moody's account with access to Moodys GenAI Ready Data.
- For **Custom OAuth app**: An OAuth app configured in your Moody's developer portal with the Amazon Quick callback URL added as a redirect URI.
- For Amazon Quick subscription requirements, see [the section called "Set up integrations in the console"](#).

Configuring Moody's

If you are using **Default OAuth app** authentication, skip this section and proceed to [the section called "Setting up the connector in Amazon Quick"](#).

For Custom OAuth app authentication, register an OAuth app in your Moody's developer portal and add the Amazon Quick callback URL `https://{region}.quicksight.aws.amazon.com/sn/oauthcallback` as a redirect URI. Replace `{region}` with your AWS Region (for example, `us-east-1`). Record the Client ID and Client Secret — you need them when you configure Amazon Quick.

Setting up the connector in Amazon Quick

Connect from the Available tab

If you want to use Default OAuth app authentication, you can connect directly from the **Available** tab without additional configuration.

1. In the Amazon Quick console, choose **Connectors**.
2. On the **Available** tab, find **Moody's GenAI Ready Data** and choose **Connect**.
3. Complete the Moody's sign-in flow and grant the requested permissions.

To configure a connector with Custom OAuth app instead, use the **Create for your team** tab as described below.

Create from the Create for your team tab

1. In the Amazon Quick console, choose **Connectors**.
2. Choose the **Create for your team** tab.
3. Find and choose **Moody's GenAI Ready Data**.
4. Enter a **Name** for your connector. Optionally, choose **+ Add Description** to add a description.
5. For **Connection type**, choose **Public network**.
6. For **OAuth Configuration**, choose one of the following authentication methods and configure the required fields.

- a. For **Default OAuth app**:

No additional credentials are needed. Choose **Next** to continue.

- b. For **Custom OAuth app**, configure the following fields:

- **Client ID** – The client ID from your Moody's OAuth app.
- **Public OAuth client** (Optional) – Select this option if your Moody's OAuth app is configured as a public client (no client secret).
- **Client secret** – The client secret from your Moody's OAuth app.
- **Token URL** – The token endpoint. Default: `https://api.moody's.com/genai-ready-data/oauth/token`
- **Authorization URL** – The authorization endpoint. Default: `https://api.moody's.com/genai-ready-data/oauth/authorize`
- **Redirect URL** – Pre-filled with the Amazon Quick callback URL.

7. Choose **Next**.
8. A Moody's authorization window opens. Review the requested permissions and choose **Allow**.
9. On the **Review** page, review the available actions for the connector. Choose **Next**.

10. On the **Publish** page, choose who can access the connector. You can enable access for everyone in your organization or search for specific teams or groups.
11. Choose **Publish**.

Available actions

After you set up the connector, the actions exposed by Moodys GenAI Ready Data are available. To see the current set of actions for your connector, go to the connector's **Available actions** view in the Amazon Quick console.

Managing and troubleshooting

To edit, share, or delete your connector, see [the section called “Managing existing integrations”](#).

Authentication issues

- **Sign-in fails (Default OAuth app or Custom OAuth app)** – Verify that your Moody's account is active and that you can sign in to the Moody's portal directly. For Custom OAuth app, confirm that the redirect URI in your Moody's OAuth app matches the Amazon Quick callback URL.
- **Invalid client credentials (Custom OAuth app)** – Verify that the Client ID and Client secret match the values in your Moody's OAuth app.

New Relic integration

With the New Relic connector, you can access the New Relic observability platform directly in Amazon Quick through natural language. You can investigate incidents, analyze application performance, query telemetry data, and generate reports without leaving Amazon Quick.

New Relic is available as a built-in connector in Amazon Quick. To set up this integration, complete the following two steps. First, prepare your New Relic account with the required access. Then, create the integration in Amazon Quick and authenticate with your New Relic credentials. This integration uses OAuth 2.0 user authentication. For information about the authentication methods that Amazon Quick supports, see [the section called “Authentication methods”](#).

Prerequisites

Before you set up the integration, make sure that you have the following resources and access.

- An active New Relic account with the required permissions and access. For information about account requirements and access configuration, see [New Relic AI MCP](#) in the New Relic documentation.
- For information about subscription requirements, see [the section called “Set up integrations in the console”](#).

Verify your New Relic account

Before you configure Amazon Quick, verify that your New Relic account meets the following requirements.

- Your New Relic account is active and you can sign in to login.newrelic.com.
- The entities and data that you want to query are accessible to your account.

Set up the integration in Amazon Quick

After you verify your New Relic account access, create the integration in Amazon Quick.

1. In the Amazon Quick console, choose **Connectors**.
2. Choose the **Create for your team** tab.
3. Find and choose **New Relic**.
4. In the **Create integration** dialog, enter the following fields:
 - **Name** – A descriptive name for your New Relic integration.
 - **Description** (Optional) – Notes about how you plan to use this connection.
 - **Connection type** – Choose **Public network**.
5. Choose **Create and continue**.
6. On the integration detail page, choose **Sign in** and authenticate with your New Relic account credentials. This step authorizes Amazon Quick to access New Relic on behalf of the authenticated user.
7. (Optional) Choose the users to share the integration with.
8. Choose **Done**.

The integration appears in the **Existing actions** panel with a status of **Available**.

Available actions

After you set up the integration, you can use the following actions.

New Relic available actions

Category	Action	Description
Incident analysis	Generate Alert Insights Report	Generates an alert intelligence analysis report for a specific issue. Provides insights into root cause, impact, and recommended actions.
Incident analysis	Generate User Impact Report	Analyzes affected entities and their relationships to quantify end-user impact through metrics like affected user count, degraded services, and severity indicators.
Incident analysis	Search Incident	Retrieves alert incident events with flexible filtering by issue ID, entity GUID, or policy and condition pair.
Incident analysis	List Recent Issues	Lists all recent issues (last 24 hours) for a specified account.
Performance analysis	Analyze Transactions	Identifies slow and error-prone transactions, their traces, common transaction paths, and error distributions.
Performance analysis	Analyze Golden Metrics	Analyzes key health indicators (throughput, response time, error rate, and saturation) across application and infrastructure entities.
Performance analysis	Analyze Deployment Impact	Compares metrics before and after a deployment to identify regressions or improvements.
Performance analysis	Analyze Kafka Metrics	Analyzes Kafka metrics including consumer lag, producer throughput, message latency, partition balance, and resource utilization.

Category	Action	Description
Performance analysis	Analyze Threads	Analyzes thread metric data including thread state, CPU usage, and memory consumption. Provides language-specific insights for applications.
Performance analysis	List Garbage Collection Metrics	Retrieves garbage collection and memory metrics for a given entity. Use this action to identify whether GC issues are affecting application performance.
Log analysis	Analyze Entity Logs	Analyzes application logs to identify error patterns, anomalous behavior, and recurring issues within a specified time window.
Log analysis	List Recent Logs	Retrieves recent logs for a specified account and entity GUID.
Error analysis	List Entity Error Groups	Retrieves error groups from the Errors Inbox within a time window. Groups errors by message and prioritizes by user impact.
Queries	Natural Language to NRQL Query	Converts a natural language request into a New Relic Query Language (NRQL) query, runs it against New Relic, and returns the results.
Queries	Execute NRQL Query	Runs an NRQL query directly against New Relic telemetry data.
Entities	Get Entity	Retrieves entities by GUID or searches by name pattern.
Entities	List Related Entities	Retrieves entities that are related to a given entity GUID.
Entities	Search Entity with Tag	Searches for entities by using tag key and value pairs.

Category	Action	Description
Entities	List Entity Types	Lists the complete catalog of New Relic entity types with their domain, type, and metric definitions.
Alerts	List Alert Policies	Lists alert policies for a specified account, with optional filtering by policy name.
Alerts	List Alert Conditions	Retrieves alert condition details for a specific alert policy.
Dashboards	List Dashboards	Lists all dashboards for a New Relic account.
Dashboards	Get Dashboard	Retrieves detailed information for a specific dashboard by entity GUID.
Change tracking	List Change Events	Retrieves a history of change events such as deployments, configuration changes, and other tracked activities for an entity.
Monitoring	List Synthetic Monitors	Lists synthetic monitors that check service availability and performance from multiple geographic locations.
Account	List Available New Relic Accounts	Lists all New Relic account IDs that are accessible to the authenticated user.
Utilities	Convert Time Period to Epoch Ms	Converts a natural language time period to epoch milliseconds for use with time-based queries.

Note

The actions that you can use depend on the permissions that are configured in your New Relic account and the entities that are accessible to the authenticated user.

Manage and troubleshoot

To edit, share, or delete your integration, see [the section called “Managing existing integrations”](#).

Authentication issues

- **Sign-in fails** – Verify that your New Relic account is active and that you can sign in to login.newrelic.com directly. If your organization uses single sign-on (SSO), confirm that your identity provider is configured correctly.
- **MCP server access denied** – Verify that your user has the required permissions to access the New Relic MCP server. For information about access requirements, see [New Relic AI MCP](#) in the New Relic documentation.

Notion integration

With Notion integration in Amazon Quick, you can manage pages, databases, and collaborative workspaces through MCP server connectivity. This integration provides action capabilities for knowledge management and content organization operations.

What you can do

Notion integration provides connector capabilities through MCP server connectivity:

- Create and edit pages and documents
- Manage databases and structured content
- Organize content with tags and properties
- Share pages and collaborate with team members
- Search across workspaces and content
- Manage workspace permissions and access

Available tools

The Notion MCP server typically provides these tools:

- `create_page` - Create new pages
- `update_page` - Update page content
- `get_page` - Retrieve page information

- `search_pages` - Search for pages
- `create_database` - Create new databases
- `query_database` - Query database entries
- `create_database_entry` - Add database entries
- `update_database_entry` - Update database entries

Note

The specific tools and capabilities available through this MCP server may change over time. For the most current information about supported tools, features, and implementation details, check the official Notion documentation and MCP server repository.

Setting up Notion integration

Notion integration uses MCP server connectivity to provide action capabilities. For detailed setup instructions, see [the section called “MCP integration”](#).

You'll need:

- Notion account with appropriate workspace permissions
- Notion integration token for API access

Compatibility

Notion integration supports:

- **Chat Agents:** Yes
- **Flows:** Yes
- **Knowledge Base:** No

OpenAPI Specification integration

With OpenAPI Specification integration, you can create custom integrations based on OpenAPI schemas. This allows you to connect to any API that provides an OpenAPI specification. For Amazon Quick subscription requirements, see [the section called “Set up integrations in the console”](#).

What you can do

OpenAPI Specification integration provides schema-based connectivity to help you work with custom APIs.

Connector

Perform actions based on OpenAPI specifications. Execute API calls, manage resources, and interact with custom services through dynamically generated actions based on the provided schema.

Schema-based configuration

Import OpenAPI specifications to automatically generate available actions and actions. Support for JSON schema validation and parameter mapping.

Before you begin

Before you set up OpenAPI Specification integration, make sure you have the following:

- OpenAPI specification (version 3.0.0 or higher) for your target API.
- API authentication credentials (API key, OAuth, or other).
- Amazon Quick Author or higher.
- API documentation and endpoint access.

Prepare OpenAPI specification and authentication

Before setting up the integration in Amazon Quick, prepare your OpenAPI specification and authentication credentials. Your OpenAPI specification must meet specific requirements for successful integration.

Basic requirements

Your OpenAPI specification must meet these basic requirements.

- **OpenAPI version** - Version 3.0.0 or higher is required.
- **File format** - JSON format only (application/json).
- **File size limit** - Maximum 1MB for the entire OpenAPI specification.
- **Operation limit** - Minimum 1 and maximum 100 API operations per connector.

Required schema fields

Your OpenAPI specification must include these required sections.

openapi

The OpenAPI version being used (must be "3.0.0" or higher).

info

Service information including title, description, and version.

servers

Base URL and server information for API connectivity.

paths

API endpoint definitions with HTTP methods and operations.

Supported authentication schemes

Prepare authentication credentials based on the supported authentication methods in OpenAPI specifications.

OAuth 2.0

Supports both Authorization Code Grant and Client Credentials Grant flows. Requires `authorizationUrl`, `tokenUrl`, and scopes definitions in the security scheme.

API Key

API key authentication passed in query parameters or headers.

No authentication

Default option when no security schemes are defined in the specification.

Note

Unsupported authentication schemes in the OpenAPI specification will be ignored, and the connector will default to no authentication.

Set up OpenAPI Specification integration

After preparing your OpenAPI specification and authentication credentials, follow these steps to create your OpenAPI Specification integration.

1. In the Amazon Quick console, choose **Connectors**.
2. Choose the **Create for your team** tab.
3. Find and choose **OpenAPI Specification**.
4. On the Add Schema page, select **Import Schema** and choose a schema to import.
5. Select **Next**.
6. Fill in the integration details including name and description.
7. Review the available actions generated from your OpenAPI specification.
8. Select **Create and continue**.
9. Choose users to share the integration with.
10. Click **Next**.

Expected results

After successful setup, your OpenAPI Specification integration appears in the integrations list with dynamically generated actions based on your schema. The integration is available for use in Amazon Quick workflows, automations, and AI agents, with tasks corresponding to the endpoints defined in your OpenAPI specification.

Format and pattern requirements

Your OpenAPI specification must follow these format requirements.

- **Path patterns** - Must follow the pattern: `^/[^\/*]*(/[^\/*]*)*$` and begin with a forward slash (/).
- **Operation IDs** - Must follow the pattern: `^[a-zA-Z0-9_]{1,256}$`.
- **Descriptions** - Required for all operations and parameters, maximum 5000 characters.
- **Content type** - Only application/json is supported for request and response bodies.

Unsupported features

The following OpenAPI features are not supported and will cause validation errors.

- **Array types** - Schemas with array types (for example, {"type": "array", "items": {"string"}}) are not supported.
- **Composition keywords** - allOf, oneOf, anyOf, and not keywords are not supported.
- **Circular references** - Circular or cyclic references (\$ref inside of \$ref) are not supported.
- **Complex nested structures** - Avoid deeply nested object structures when possible.

Schema best practices

Follow these best practices when creating your OpenAPI specification.

Write effective descriptions

Use these guidelines to write clear and helpful descriptions.

- **Operation descriptions** - Describe what the operation does, when to use it, and how it behaves.
- **Parameter descriptions** - Concisely explain the parameter's purpose and impact on operation behavior.
- **Self-contained content** - Ensure descriptions provide sufficient guidance without external references.
- **Dependency clarity** - Make dependencies between operations explicit in descriptions.
- **Operation references** - Use operationId when referring to other operations, not API paths.

Parameter structure recommendations

Structure your parameters for optimal usability.

- **Flatten complex objects** - Instead of nested objects, use separate parameters (for example, start_date, start_time, end_date, end_time).
- **Use standard formats** - Use standard format values like "date-time" or "date" for ISO-8601 dates and times.
- **Clear parameter names** - Use descriptive parameter names that clearly indicate their purpose.
- **Required field marking** - Properly mark parameters as required or optional based on API behavior.

Supported extensions

We support custom OpenAPI extensions to enhance the user experience.

x-amzn-form-display-name

Use at the parameter level to override the default field name displayed in confirmation forms. Currently supported for request body parameters only.

```
"x-amzn-form-display-name": "User Name"
```

x-amzn-operation-type

Define whether an operation should be treated as "read" or "write". By default, GET methods are "read" operations and all other HTTP methods are "write" operations.

```
"x-amzn-operation-type": "read"
```

Schema validation and error handling

When you upload an OpenAPI specification, we perform comprehensive validation.

- **File size validation** - Ensures the specification doesn't exceed 1MB.
- **Operation count validation** - Verifies between 1-100 operations are defined.
- **Schema structure validation** - Checks for required fields and proper formatting.
- **Security scheme validation** - Validates supported authentication methods.
- **Content type validation** - Ensures only application/json is used.

Validation errors appear below the schema editor and must be resolved before the integration can be created. Common validation issues include:

- Missing required fields (openapi, info, servers, paths).
- Invalid path patterns or operation IDs.
- Unsupported schema features (arrays, composition keywords).
- Missing or overly long descriptions.

- Circular references in schema definitions.

Manage OpenAPI Specification integration

After you create your OpenAPI Specification integration, you can manage it through several options.

Share integration

You can share OpenAPI Specification connectors with other users in your organization.

1. From the OpenAPI integration details page, choose **Share**.
2. Configure sharing options:
 - **Share with specific users** - Enter user names or email addresses.
 - **Share with organization** - Make available to all users in your organization.
3. Set permission levels for shared access.
4. Choose **Share integration** to apply the sharing settings.

Delete integration

Follow these steps to permanently remove your OpenAPI integration.

1. From the OpenAPI integration details page, choose **Delete**.
2. Review the deletion impact, including any workflows or automations using this integration.
3. Type the integration name to confirm deletion.
4. Choose **Delete integration** to permanently remove it.

Troubleshoot OpenAPI Specification integration

Use these troubleshooting tips to resolve common OpenAPI Specification integration issues.

Schema validation errors

Ensure your OpenAPI specification follows the correct format and includes all required fields. Use an OpenAPI validator to check your schema before importing.

No tasks generated

Verify that your OpenAPI specification includes path definitions with HTTP methods and operation IDs. Actions are generated based on the operations defined in your schema.

Authentication failures

Check that the authentication scheme defined in your OpenAPI specification matches the actual API requirements and that credentials are properly configured.

Action execution failures

Review the action parameters and ensure they match the parameter definitions in your OpenAPI specification, including required fields and data types.

PagerDuty integration

Connect Amazon Quick to your PagerDuty system to manage incidents, alerts, schedules, and on-call rotations. You can perform these actions without leaving your Amazon Quick environment. For Amazon Quick subscription requirements, see [the section called "Set up integrations in the console"](#).

What you can do

With PagerDuty integration, you can perform actions within your PagerDuty systems through the PagerDuty API.

Connector

Create, update, and manage incidents, alerts, schedules, and escalation policies through the PagerDuty API.

Set up PagerDuty integration

Follow these steps to connect Amazon Quick to your PagerDuty system.

1. In the Amazon Quick console, choose **Connectors**.
2. Choose the **Create for your team** tab.
3. Find and choose **PagerDuty**.
4. Fill in the integration details:

- **Name** - Enter a descriptive name for your PagerDuty integration.
 - **Description** (Optional) - Describe the purpose of this integration.
5. Choose your connection type:
 - **User authentication** - OAuth for individual user access to PagerDuty.
 - **Service authentication** - API key-based authentication for service-to-service access.
 6. Fill in the connection settings based on your selected authentication method (either user or service).
 7. Select **Create and continue**.
 8. Choose users to share the integration with.
 9. Click **Next**.

What happens next

After you complete the setup, your PagerDuty integration appears in the integrations list. You can use it in Amazon Quick workflows, automations, and AI agents.

Manage PagerDuty integrations

You can perform these management tasks for your PagerDuty integrations:

- **Edit integration** - Update authentication settings or PagerDuty configuration.
- **Share integration** - Make the integration available to other users in your organization.
- **Delete integration** - Remove the integration and revoke authentication.

PagerDuty Advance integration

Connect Amazon Quick to the PagerDuty Advance MCP and use AI to securely contact PagerDuty Advance. Teams can access incident context, history, and runbooks from the SRE agent, query the Insights agent for proactive ways to improve, or alter on-call coverage with Shift agent from their AI-enabled platform of choice.

What you can do

The PagerDuty Advance integration enables you to:

- Access incident context, history, and runbooks for active incident response
- Perform root cause analysis and technical troubleshooting
- Query on-call schedules and manage shift coverage
- Analyze historical trends and performance metrics (MTTR, MTTA)
- Generate insights for service reliability and team performance
- Collaborate with PagerDuty Advance agents from your AI-enabled platform of choice

Available tools

The MCP server for PagerDuty Advance typically provides these tools:

- `analytics_agent_tool` - Runs an analytics agent for historical reporting and trend analysis, including performance metrics (MTTR, MTTA), team benchmarking, service reliability statistics, and executive KPI tracking.
- `schedules_agent_tool` - Runs a schedules agent for on-call and scheduling information, including shift assignments, schedule configurations, escalation policies, user availability conflicts, and coverage management.
- `sre_agent_tool` - Runs an SRE agent for incident response and technical troubleshooting, including active incident triage, root cause analysis, diagnostics, alert explanations, change event analysis, and runbook generation.

Note

The specific tools and capabilities available through this MCP server may change over time. For the most current information about supported tools, features, and implementation details, check the official PagerDuty documentation and MCP server repository.

Setting up PagerDuty Advance integration

PagerDuty Advance integration uses MCP server connectivity to provide agent capabilities. For detailed setup instructions, see [the section called “MCP integration”](#).

You'll need:

- PagerDuty account with appropriate permissions

- PagerDuty API credentials (Client ID and Client Secret)
- Amazon Quick with MCP integration enabled

Compatibility

PagerDuty Advance integration supports:

- **Chat Agents:** Yes
- **Flows:** Yes
- **Knowledge Base:** No

QuickBooks integration

With the QuickBooks connector, you can access QuickBooks Online directly in Amazon Quick through natural language. You can manage invoices, customers, vendors, accounts, and generate financial reports without leaving Amazon Quick.

Amazon Quick supports multiple authentication methods for QuickBooks. Choose the method that best fits your organization's security requirements.

- **Default OAuth app** – Uses an AWS-managed OAuth application. No additional credentials are needed. Users authenticate directly with their Intuit account.
- **Custom OAuth app** – Uses a customer-managed application registered in the Intuit Developer portal. This option gives your organization full control over the OAuth configuration.

For more information about the authentication methods that Amazon Quick supports, see [the section called “Authentication methods”](#).

Before you begin

Make sure that you have the following before you set up the integration.

- An active QuickBooks Online account with access to the company data that you want to use.
- For **Custom OAuth app**: Access to the [Intuit Developer portal](#) on the Intuit website to create an app.
- For Amazon Quick subscription requirements, see [the section called “Set up integrations in the console”](#).

Configuring QuickBooks

If you are using **Default OAuth app** authentication, skip this section and proceed to [the section called "Setting up the connector in Amazon Quick"](#).

For Custom OAuth app authentication, complete the following steps in the Intuit Developer portal before configuring Amazon Quick.

Create an Intuit Developer app (Custom OAuth app)

Create an app in the Intuit Developer portal to obtain the client credentials that you need for Amazon Quick. For more information, see [OAuth 2.0](#) on the Intuit Developer website.

1. Sign in to the [Intuit Developer portal](#) on the Intuit website.
2. From your workspace, choose the plus (+) button to create a new app.
3. For **Type**, select the QuickBooks Online platform.
4. Enter an **App name** for your application.
5. For **Permissions**, under **QuickBooks Online and Payments**, select the authorization scopes that your integration requires. For the recommended scopes, see [the section called "Recommended scopes"](#).
6. Complete the app creation. On the confirmation page, choose **Show credentials** to view your development credentials.
7. Record the following values. You need them when you configure Amazon Quick.
 - **Client ID**
 - **Client secret**
8. In the app settings, choose the **Redirect URIs** tab. Add the Amazon Quick callback URL:
`https://{region}.quicksight.aws.amazon.com/sn/oauthcallback`
9. Choose **Save**.

Recommended scopes

When you create an Intuit Developer app, select the following authorization scopes based on the actions that you want to use.

QuickBooks recommended scopes

Scope	Description
com.intuit.quickbooks.accounting	Accesses QuickBooks Online accounting data, including invoices, customers, vendors, accounts, and financial reports.
com.intuit.quickbooks.payment	Accesses QuickBooks Online payment data, including purchases and payment transactions.
openid	Authenticates the user's identity.
email	Reads the user's email address.
profile	Reads the user's profile information.
phone	Reads the user's phone number.
address	Reads the user's address.

Note

The openid, email, profile, phone, and address scopes are used for user authentication with Custom OAuth app.

Setting up the connector in Amazon Quick

Connect from the Available tab

If you want to use Default OAuth app authentication, you can connect directly from the **Available** tab without additional configuration.

1. In the Amazon Quick console, choose **Connectors**.
2. On the **Available** tab, find **QuickBooks** and choose **Connect**.
3. Complete the QuickBooks sign-in flow and grant the requested permissions.

To configure a connector with Custom OAuth app instead, use the **Create for your team** tab as described below.

Create from the Create for your team tab

After you complete any required QuickBooks configuration, create the connector in Amazon Quick.

1. In the Amazon Quick console, choose **Connectors**.
2. Choose the **Create for your team** tab.
3. Find and choose **QuickBooks**.

Note

If a QuickBooks connector already exists, a dialog appears with your existing connectors. To use an existing connector, choose it. To create a new one, choose **No, create new**.

4. Enter a **Name** for your connector. Optionally, choose **+ Add Description** to add a description.
5. For **Connection type**, choose **Public network**.
6. For **OAuth Configuration**, choose one of the following authentication methods and configure the required fields.
 - a. For **Default OAuth app**:

No additional credentials are needed. Choose **Next** to continue.
 - b. For **Custom OAuth app**, configure the following fields:
 - **Base URL** (Optional) – The QuickBooks API base URL. Example: `https://quickbooks.api.intuit.com`
 - **Client ID** – The client ID from your Intuit Developer app.
 - **Client secret** – The client secret from your Intuit Developer app.
 - **Token URL** – The token endpoint. Example: `https://oauth.platform.intuit.com/oauth2/v1/tokens/bearer`
 - **Authorization URL** – The authorization endpoint. Example: `https://appcenter.intuit.com/connect/oauth2`
 - **Redirect URL** – Pre-filled with the Amazon Quick callback URL.
7. Choose **Next**.

8. If you chose **Default OAuth app** or **Custom OAuth app**, an Intuit authorization window opens. Review the requested permissions and choose **Connect**.
9. On the **Review** page, review the available actions for the connector. Choose **Next**.
10. On the **Publish** page, choose who can access the connector. You can enable access for everyone in your organization or search for specific teams or groups.
11. Choose **Publish**.

Available actions

After you set up the connector, the following actions are available.

QuickBooks available actions

Category	Action	Description
Queries	Search Entities	Searches across QuickBooks entities using a query.
Accounts	Get Account Details	Retrieves details for a specific account by query.
Accounts	Get Account	Retrieves a specific account by ID.
Accounts	Create Account	Creates a new account in the chart of accounts.
Invoices	Get Invoice	Retrieves a specific invoice by ID.
Invoices	List Invoices	Lists invoices with optional filtering.
Invoices	Create Invoice	Creates a new invoice.
Customers	Get Customer	Retrieves a specific customer by ID.
Customers	Create Customer	Creates a new customer record.
Vendors	Create Vendor	Creates a new vendor record.
Purchases	Create Purchase	Creates a new purchase transaction.

Category	Action	Description
Purchases	Create Bill	Creates a new bill from a vendor.
Journal entries	Create Journal Entry	Creates a new journal entry.
Reports	Get Profit and Loss Report	Generates a profit and loss report.
Reports	Get Customer Balance Report	Generates a customer balance summary report.
Reports	Get Customer Balance Detail	Generates a detailed customer balance report.
Reports	Get Vendor Balance Report	Generates a vendor balance summary report.
Reports	Get Vendor Balance Detail	Generates a detailed vendor balance report.
Reports	Get General Ledger Report	Generates a general ledger report.
Batch operations	Batch Execute Operations	Runs multiple operations in a single batch request.
Company	Get Company Info	Retrieves information about the connected company.

Note

The actions that you can use depend on the company data accessible to the authenticated user.

Managing and troubleshooting

To edit, share, or delete your connector, see [the section called “Managing existing integrations”](#).

Authentication issues

- **Sign-in fails (Default OAuth app or Custom OAuth app)** – Verify that your Intuit account is active and that you can sign in to quickbooks.intuit.com on the Intuit website directly. For

Custom OAuth app, confirm that the redirect URI in your Intuit Developer app matches the Amazon Quick callback URL.

- **Invalid client credentials (Custom OAuth app)** – Verify that the Client ID and Client secret match the values in your Intuit Developer app. You can view your credentials from the app settings in the Intuit Developer portal.
- **Insufficient permissions** – Verify that the scopes configured for your Intuit Developer app include the permissions required for the actions that you want to use. See [the section called “Recommended scopes”](#).

REST API Connection integration

With REST API Connection integration in Amazon Quick, you can perform actions with custom REST APIs and web services. This integration supports action execution only.

What you can do

With REST API Connection integration, you can perform actions with custom REST APIs and web services through the connector.

Connector

Perform HTTP requests, retrieve data, and interact with APIs using flexible authentication options.

Note

REST API Connection integration doesn't support data access or knowledge base creation. It's designed specifically for task execution and API interactions with custom web services.

Before you begin

Before you set up REST API integration, make sure you have the following:

- REST API endpoint with appropriate access permissions.
- API authentication credentials (OAuth, API key, or other).
- API documentation for the target web service.

Prepare API endpoint and authentication

Before configuring the integration in Amazon Quick, prepare your REST API endpoint and authentication credentials. REST API Connection integration supports multiple authentication methods. Choose the method that matches your API requirements:

User authentication (OAuth)

Gather the following information from your API provider:

- **Base URL** - REST API base URL.
- **Client ID** - OAuth application client ID.
- **Client Secret** - OAuth application client secret.
- **Token URL** - OAuth token endpoint.
- **Auth URL** - OAuth authorization endpoint.
- **Redirect URL** - OAuth redirect URI.

Service authentication (Service-to-service OAuth)

Gather the following information from your API provider:

- **Authentication type** - OAuth 2.0 client credentials grant flow for service-to-service authentication.
- **Base URL** - REST API base URL.
- **Client ID** - OAuth application client identifier for service authentication.
- **Client secret** - OAuth application client secret for service authentication.
- **Token URL** - OAuth token endpoint for obtaining access tokens.

Custom headers and parameters

You can use custom headers and parameters for flexible authentication and API interaction:

- Custom authentication headers.
- API version headers.
- Content-Type specifications.
- Custom query parameters.

Set up REST API integration

After you have prepared your API endpoint and authentication credentials, follow these steps to create your REST API integration:

1. In the Amazon Quick console, choose **Connectors**.
2. Choose the **Create for your team** tab.
3. Find and choose **REST API Connection**.
4. Fill in the integration details:
 - **Name** - Descriptive name for your REST API integration.
 - **Description** (Optional) - Purpose of the integration.
5. Choose your connection type:
 - **User authentication** - OAuth-based authentication for individual user access.
 - **Service authentication** - API key-based authentication for service access.
6. Fill in the connection settings based on your selected authentication method (either user or service).
7. Select **Next**.
8. Review the actions that are available.
9. Select **Create and continue**.

Available task actions

After you create your REST API integration, you can review the available actions for interacting with the REST API. Common REST API actions include:

- HTTP GET requests for data retrieval.
- HTTP POST requests for data creation.
- HTTP PUT/PATCH requests for data updates.
- HTTP DELETE requests for data removal.
- Custom endpoint interactions.
- JSON and XML data processing.
- Query parameter and header management.

 Note

The specific actions available depend on the REST API endpoints and the authentication permissions configured for your integration.

API configuration options

You can configure various aspects of your REST API integration to match your specific requirements.

Endpoint configuration

Configure these endpoint settings:

- Base URL and endpoint paths.
- HTTP method specifications.
- Request and response format handling.
- Error handling and retry logic.

Data handling

Configure how your integration processes different data formats:

- JSON request and response processing.
- XML data transformation.
- Form data and multipart uploads.
- Binary data handling.

Manage REST API integrations

After you create your REST API integration, you can manage it using these options:

- **Edit integration** - Update authentication settings, base URL, or API configuration.
- **Share integration** - Make the integration available to other users in your organization.
- **Monitor usage** - View integration activity and API call metrics.

- **Review actions** - See the complete list of available REST API actions.
- **Test endpoints** - Validate API connectivity and authentication.
- **Delete integration** - Remove the integration and revoke associated authentication.

Important

REST API integrations depend on the availability and configuration of the target web service. Changes to the API or authentication requirements may affect integration functionality.

Salesforce integration

With the Salesforce connector in Amazon Quick, you can perform actions within Salesforce organizations, including managing records, querying data, and interacting with Salesforce APIs. For Amazon Quick subscription requirements, see [the section called "Set up integrations in the console"](#).

Amazon Quick supports the following authentication options for Salesforce.

Salesforce authentication flows

Flow	Description
Managed OAuth (Default OAuth app)	Uses an AWS managed OAuth application with no additional credentials or External Client App configuration required. Requires the Amazon Quick Salesforce package and an activated Model Context Protocol (MCP) server in Salesforce.
User authentication (three-legged OAuth, 3LO)	Authorization code flow. Each user authenticates through a browser. Best for multi-user environments where actions run on behalf of the signed-in user.
Service authentication (two-legged OAuth, 2LO)	Client credentials flow. Machine-to-machine authentication with no user interaction. All actions run as a designated "Run As" user that you configure in Salesforce.

Note

As of the Salesforce Spring '26 release (February 2026), new Connected App creation is disabled by default across all Salesforce organizations. This guide uses Salesforce's External Client Apps (ECAs), which are fully compatible with the Amazon Quick Salesforce connector.

Prerequisites

Before you set up the Salesforce integration, make sure you have the following:

- A Salesforce organization with System Administrator access.
- Amazon Quick Author role or higher.
- Your Amazon Quick instance URL (for example, `https://us-east-1.quicksight.aws.amazon.com`).
- For Managed OAuth: The [Amazon Quick Salesforce package](#) installed from Salesforce AppExchange.
- The headless-360 MCP server activated in your Salesforce organization. This is required for all authentication flows. For activation steps, see [the section called "Activate the MCP server"](#).

Activate the MCP server

Before you configure any authentication flow, activate the headless-360 MCP server in your Salesforce organization. This step is required for Managed OAuth, user authentication (3LO), and service authentication (2LO).

1. Sign in to your Salesforce organization.
2. From **Setup**, in the **Quick Find** box, enter MCP Servers, and then select **MCP Servers**.
3. Find headless-360 in the list of servers.
4. Choose **Activate**.

Set up Managed OAuth

If you are using user authentication (three-legged OAuth, 3LO) or service authentication (two-legged OAuth, 2LO), skip this section and proceed to [the section called "Step 1: Create an External Client App in Salesforce"](#).

Managed OAuth is the simplest setup option. It uses an AWS managed OAuth application, so you do not need to create an External Client App or configure OAuth credentials manually.

Install the Amazon Quick Salesforce package

1. Sign in to your Salesforce organization as a System Administrator.
2. Go to the Amazon Quick Salesforce package listing on Salesforce AppExchange (linked in the prerequisites).
3. Choose **Get It Now** and follow the installation prompts to install the package in your Salesforce organization.

After you install the package, proceed to [the section called "Step 4: Configure Salesforce connector in Amazon Quick"](#) to configure the connector in Amazon Quick. In Step 4, choose **Default OAuth app** as the **Auth Type** instead of **Custom OAuth app**.

Step 1: Create an External Client App in Salesforce

Create an External Client App (ECA) in Salesforce to enable OAuth authentication with Amazon Quick.

Navigate to External Client App Manager

1. Sign in to your Salesforce organization and choose **Setup** (gear icon in the top right).
2. In the navigation pane, expand **Platform Tools > Apps > External Client Apps > External Client App Manager**.
3. Choose **New External Client App**.

Fill in basic information

Enter the following basic information for your External Client App:

External Client App basic information

Field	Value
External Client App Name	A descriptive name (for example, "Amazon Quick Integration")
API Name	Auto-generated from the app name

Field	Value
Contact Email	Your admin email address
Distribution State	Local (default)

Step 2: Configure OAuth settings

Enable OAuth

In the **API (Enable OAuth Settings)** section, select **Enable OAuth** to expand the OAuth configuration panel.

Set callback URL

Enter the Amazon Quick OAuth callback URL in the **Callback URL** field:

<your-quick-instance-url>/sn/oauthcallback

For example: `https://us-east-1.quicksight.aws.amazon.com/sn/oauthcallback`

Add OAuth scopes

Move the following scopes from **Available OAuth Scopes** to **Selected OAuth Scopes**:

- Access the identity URL service (id, profile, email, address, phone)
- Manage user data via APIs (api)
- Manage user data via Web browsers (web)
- Full access (full)
- Access Visualforce applications (visualforce)
- Perform requests at any time (refresh_token, offline_access)
- Access unique user identifiers (openid)
- Access custom permissions (custom_permissions)
- Access Connect REST API resources (chatter_api)
- Access Analytics REST API resources (wave_api)
- Access Analytics REST API Charts Geodata resources (eclair_api)
- Manage Pardot services (pardot_api)

Enable OAuth flows

Select the following flow options:

For both user authentication (3LO) and service authentication (2LO):

- **Enable Client Credentials Flow**
- **Enable Authorization Code and Credentials Flow**
- **Enable Token Exchange Flow**

For user authentication (3LO) only:

- **Require user credentials in the POST body for Authorization Code and Credentials Flow.** This setting sends user credentials in the POST body during authorization code exchange. It applies only to user authentication and has no effect on service authentication.

Note

If you plan to use only service authentication (2LO), you do not need to select **Require user credentials in the POST body for Authorization Code and Credentials Flow**.

Enable security settings

Select the following security options:

For user authentication (3LO):

- **Require Proof Key for Code Exchange (PKCE).** Select this option.
- **Require Secret for Web Server Flow** (selected by default)
- **Require Secret for Refresh Token Flow** (selected by default)

For service authentication (2LO):

- **Require Proof Key for Code Exchange (PKCE).** Clear this option. PKCE is designed for browser-based authorization code flows and is incompatible with client credentials.
- **Require Secret for Web Server Flow** (selected by default)

- **Require Secret for Refresh Token Flow** (selected by default)

Important

If you need to support both user authentication (3LO) and service authentication (2LO) on the same External Client App, clear the **Require Proof Key for Code Exchange (PKCE)** option. PKCE cannot be enabled alongside client credentials on the same app. Create separate External Client Apps if you require PKCE for user authentication.

Choose **Create** to save the External Client App.

Step 3: Get consumer credentials

1. After creation, you are on the ECA detail page.
2. Navigate to the **Settings** tab or choose **Manage Consumer Details**.
3. You might need to verify your identity (email verification code).
4. Copy the **Consumer Key** (Client ID) and **Consumer Secret** (Client Secret).

Important

Save these credentials securely. You need them when you configure the connector in Amazon Quick.

Step 3b: Configure service authentication (2LO)

If you are using user authentication (3LO) only, skip this section and proceed to [the section called "Step 4: Configure Salesforce connector in Amazon Quick"](#).

Service authentication requires additional configuration in the External Client App's Policies tab. These steps enable the client credentials flow operationally and designate the user identity under which all actions run.

Enable client credentials and set the Run As user

1. In Salesforce Setup, navigate to your External Client App and choose the **Policies** tab.

2. Choose **Edit**.
3. Under **OAuth Policies**, select **Enable Client Credentials Flow**.
4. In the **Run As** field that appears, choose the user account under which all service authentication actions run. This is typically an administrator or dedicated service account.
5. Choose **Save**.

Important

The **Enable Client Credentials Flow** option exists in two places: the **Settings** tab (enables the capability) and the **Policies** tab (enables it operationally and sets the Run As user). You must enable it in both places. Salesforce tokens are always user-scoped. The Run As user determines the identity and permissions for all actions that run through service authentication.

Verify the token endpoint

Before you configure Amazon Quick, verify that the client credentials flow works by requesting a token from the Salesforce token endpoint. Run the following command from a terminal:

```
curl -X POST https://your-domain.my.salesforce.com/services/oauth2/token \  
-d "grant_type=client_credentials" \  
-d "client_id=CONSUMER_KEY" \  
-d "client_secret=CONSUMER_SECRET"
```

A successful response returns a JSON object that contains `access_token`, `instance_url`, and `token_type` fields.

If the response contains the error `no_client_credentials_user_enabled`, the Run As user is not configured. Return to [the section called “Enable client credentials and set the Run As user”](#) and verify that you enabled client credentials and assigned a Run As user on the **Policies** tab.

Step 4: Configure Salesforce connector in Amazon Quick

Navigate to Connectors

1. Open Amazon Quick and choose **Connectors**.
2. Choose the **Create for your team** tab.

3. Find and choose **Salesforce**.
4. If prompted that a connector already exists, choose **No, create new**.

Managed OAuth connection details

If you are using Managed OAuth, enter the following connection details.

Name

A descriptive name (for example, "Salesforce Managed").

Network

Public network.

Auth Type

Default OAuth app.

No other fields are required. You do not need to enter a Client ID, Client Secret, Token URL, or Authorization URL.

Choose **Next**. Complete the Salesforce sign-in flow and grant the requested permissions.

User authentication (3LO) connection details

If you are using user authentication, enter the following connection details:

Salesforce connector: user authentication (3LO)

Field	Value
Name	A descriptive name (for example, "Salesforce ECA")
Network	Public network
Auth Type	Custom OAuth app
Base URL	https://<your-domain>.my.salesforce.com/services/data/v60.0
Client ID	Consumer Key from Step 3

Field	Value
Client Secret	Consumer Secret from Step 3
Token URL	https://<your-domain>.my.salesforce.com/services/oauth2/token
Authorization URL	https://<your-domain>.my.salesforce.com/services/oauth2/authorize
Redirect URL	Pre-filled (do not change)

Choose **Next** to proceed.

Service authentication (2LO) connection details

If you are using service authentication, enter the following connection details:

Salesforce connector: service authentication (2LO)

Field	Value
Name	A descriptive name (for example, "Salesforce Service Auth")
Network	Public network
Auth Type	Custom OAuth (Service Auth)
Base URL	https://<your-domain>.my.salesforce.com/services/data/v60.0
Client ID	Consumer Key from Step 3
Client Secret	Consumer Secret from Step 3
Token URL	https://<your-domain>.my.salesforce.com/services/oauth2/token
Authorization URL	Leave blank. Do not enter a value for service authentication.
Scopes	api (use space-separated short-form scope names only)

⚠ Important

Do not populate **Authorization URL** for service authentication. Amazon Quick rejects the form if an authorization URL is entered for a client credentials flow.

Choose **Next** to proceed.

Step 5: Verify and publish

1. On the **Review** page, verify the available Salesforce actions (up to 42 actions).
2. Choose **Next** to proceed to user sharing.
3. Select users or groups who should have access to this connector.
4. Choose **Publish** to publish the connector.

After you publish the connector, you can use Salesforce actions in Amazon Quick chat, flows, and automations.

Available actions

After you set up the connector, the Salesforce Headless 360 MCP server provides four tools that Amazon Quick uses to interact with Salesforce.

Discover

Finds available Salesforce operations by running a semantic search across the operation index. Returns a ranked set of candidate operations that match your request.

Describe

Returns the technical contract for an operation, including APIs, parameters, dependencies, and ordered steps.

Dispatch

Runs the chosen operation. Routes the request to the correct endpoint and enforces access controls before the operation runs.

Dispatch (Read-Only)

Runs read-only operations. This tool never changes data or configuration. Use it when you only need to retrieve information.

Through these tools, you can perform tasks such as the following.

- Query, create, and update Salesforce records.
- Manage users, including creating, deactivating, freezing, and assigning permission sets.
- Read, write, and deploy Apex triggers.
- Build event-driven integrations with platform events, Change Data Capture, and event relays.
- Create named credentials, including authentication mechanisms and endpoints.
- Manage Commerce Cloud orders.

 Note

The library of available operations grows with each Salesforce release. Your connector discovers the current set of operations at runtime through the Discover tool.

 Note

Salesforce integration supports action execution only. Data access and knowledge base creation are not available for Salesforce systems.

Troubleshooting

To edit, share, or delete your connector, see [the section called “Managing existing integrations”](#).

- **URL no longer exists error.** Verify that you are using the correct Salesforce domain in the Token URL and Authorization URL. Do not use legacy endpoints.
- **OAuth validation fails (user authentication).** Verify that PKCE is enabled on the ECA and that the Callback URL exactly matches your Amazon Quick instance URL with `/sn/oauthcallback` appended.
- **Missing scopes error.** Go back to the ECA configuration in Salesforce and ensure all required OAuth scopes are in the **Selected** list, not still in **Available**.
- **Connected App option is unavailable.** This is expected after Salesforce Spring '26. Use an External Client App instead. ECAs are the supported replacement.

Service authentication (2LO) issues

- **no client credentials user enabled error.** The client credentials flow is not enabled on the **Policies** tab, or no Run As user is assigned. Navigate to the External Client App, choose the **Policies** tab, choose **Edit**, select **Enable Client Credentials Flow**, assign a Run As user, and save. For more information, see [the section called “Enable client credentials and set the Run As user”](#).
- **One or more parameters are invalid error in Amazon Quick.** This error typically indicates one of the following issues:
 - The **Authorization URL** field contains a value. Leave this field blank for service authentication.
 - The **Scopes** field uses commas instead of spaces. Use space-separated short-form scope names only.
 - The **Client ID** or **Client Secret** contains leading or trailing whitespace.
 - The **Token URL** contains a trailing slash.
- **request not supported on this domain error.** Use your Salesforce runtime domain (*your-domain*.my.salesforce.com), not the setup domain (*.salesforce-setup.com).
- **PKCE errors with service authentication.** PKCE is incompatible with the client credentials flow. Clear the **Require Proof Key for Code Exchange (PKCE)** option in the External Client App's security settings. For more information, see [the section called “Enable security settings”](#).

SAP workload integrations

With SAP workload integrations, you can perform actions within various SAP systems. Manage business data, inventory, materials, and business processes. For Amazon Quick subscription requirements, see [the section called “Set up integrations in the console”](#).

What you can do

SAP workload integrations provide enterprise-level connectivity to help you work with your SAP systems.

Connector

Perform actions within SAP systems. Create, update, and manage business data, inventory records, material information, and other enterprise operations through SAP APIs.

Multiple SAP modules

Support for five different integration types: Bill of Materials, Business Partner, Material Stock, Physical Inventory Documents, and Product Master.

Before you begin

Before you set up SAP workload integrations, make sure you have the following:

- SAP system with appropriate modules installed and configured.
- SAP user account with necessary permissions and authorizations.
- Amazon Quick Author or higher.
- Network connectivity between Amazon Quick and your SAP system.

Prepare SAP system configuration and authentication

Before setting up the integration in Amazon Quick, prepare your SAP system configuration and user authentication. SAP workload integrations support multiple authentication methods and require proper system setup.

Authentication methods

SAP workload integrations support two authentication methods:

OAuth 2.0 (Recommended)

Secure authentication method for automated workflows. Requires OAuth configuration in your SAP system.

Required parameters:

- **Client ID** - SAP OAuth Client ID
- **Client Secret** - SAP OAuth Client Secret
- **Token URL** - OAuth token endpoint (e.g., `https://hostname:port/sap/bc/sec/oauth2/token?sap-client=100`)
- **Domain URL** - SAP system API endpoint (e.g., `https://hostname:port/sap/opu/odata/sap/API_BUSINESS_PARTNER`)

Basic Authentication

Username and password authentication for direct SAP system access.

Required parameters:

- **Username** - SAP system username
- **Password** - SAP system password
- **Domain URL** - SAP system API endpoint

SAP system configuration requirements

Before using SAP workload integrations, ensure your SAP system is properly configured:

OAuth 2.0 setup

For OAuth 2.0 authentication:

- Configure the OAuth authorization server with appropriate scopes
- Verify OAuth is enabled for the service using transaction code /IWFND/MAINT_SERVICE
- Refer to SAP documentation for detailed OAuth 2.0 configuration: [OAuth 2.0 Configuration Guide](#)

SAP API activation

Ensure the required SAP API services are active:

- Activate the specific API service for your chosen SAP connector
- Verify API service status in your SAP system
- Refer to SAP documentation for API activation: [SAP API Service Activation Guide](#)

Available SAP workload integration types

Choose the SAP integration type that matches your business needs and ensure the corresponding SAP module is available in your system.

SAP Bill of Materials

Manage bill of materials data, including component lists, quantities, and manufacturing specifications.

Required scope: ZAPI_BILL_OF_MATERIAL_SRV_0002

SAP Business Partner

Handle business partner information, including customer and vendor data, contact details, and relationship management.

Required scope: ZAPI_BUSINESS_PARTNER_0001

SAP Material Stock

Access and manage material stock levels, inventory movements, and warehouse information.

Required scope: ZAPI_MATERIAL_STOCK_SRV_0001

SAP Physical Inventory Documents

Create and manage physical inventory documents, stock counts, and inventory reconciliation processes.

Required scope: ZAPI_MATERIAL_STOCK_SRV_0001

SAP Product Master

Maintain product master data, including material specifications, classifications, and product hierarchies.

Required scope: ZAPI_PRODUCT_SRV_0001

SAP workload authentication setup

Prepare SAP user authentication credentials and ensure proper system access permissions.

- **SAP system URL** - Obtain the base URL or server address for your SAP system.
- **User credentials** - Create or identify a SAP user account with appropriate permissions.
- **System permissions** - Ensure the user account has the necessary SAP authorizations and transaction codes for the specific SAP modules you plan to integrate.
- **Network access** - Verify that your SAP system is accessible from external connections and that firewall rules allow the integration.

Set up SAP integrations

After preparing your SAP system configuration and authentication credentials, the setup process is similar for all SAP integration types. Choose between OAuth 2.0 or Basic Authentication based on your security requirements.

1. In the Amazon Quick console, choose **Connectors**.
2. Choose the **Create for your team** tab.
3. Select one of the available SAP connectors:
 - **SAP Bill of Materials**
 - **SAP Business Partner**
 - **SAP Material Stock**
 - **SAP Physical Inventory Documents**
 - **SAP Product Master**
4. Click **Add** (plus "+" button).
5. Fill in the name and description for your SAP integration.
6. Choose the connection type for your integration.
7. Configure authentication using one of the supported methods:

- a. **OAuth 2.0 configuration**

For OAuth 2.0 authentication, provide:

- **Client ID** - Your SAP OAuth Client ID
- **Client Secret** - Your SAP OAuth Client Secret
- **Token URL** - OAuth token endpoint URL
- **Domain URL** - SAP system API endpoint URL

- b. **Basic Authentication configuration**

For Basic Authentication, provide:

- **Username** - Your SAP system username
- **Password** - Your SAP system password
- **Domain URL** - SAP system API endpoint URL

8. Select **Create and continue**.

9. Choose users to share the integration with.
10. Click **Next**.

Expected results

After successful setup, your SAP workload integration appears in the integrations list and is available for use in Amazon Quick workflows, automations, and AI agents. You can perform SAP-specific actions directly from Amazon Quick using the configured authentication credentials.

Available operations by connector type

Each SAP workload connector provides specific operations tailored to its business domain. Review the available operations for your chosen connector type.

SAP Bill of Materials operations

Available operations for managing bill of materials data:

- **Get Material BOM Item** - Retrieves bill of material details for specified materials

Operation ID: getMaterialBOMItem

Endpoint: GET /MaterialBOMItem

SAP Business Partner operations

Available operations for managing business partner data:

- **Get Business Partner** - Retrieves business partner general data

Operation ID: getBusinessPartner

- **Get Business Partner Addresses** - Retrieves business partner address data

Operation ID: getBusinessPartnerAddress

- **Get Business Partner Roles** - Retrieves business partner role data

Operation ID: getBusinessPartnerRole

- **Get Business Partner by ID** - Retrieves business partner data by business partner number

Operation ID: getBusinessPartnerByID

Required parameter: BusinessPartner (string, max 10 characters)

- **Get Business Partner Role by ID** - Retrieves business partner role data using key fields

Operation ID: getBusinessPartnerRoleByID

Required parameters: BusinessPartner (string, max 10 characters), BusinessPartnerRole (string, max 6 characters)

SAP Material Stock operations

Available operations for managing material stock data:

- **Get Material Stock In Account** - Retrieves material stock information posted in the account model

Operation ID: getMaterialStockInAccount

SAP Physical Inventory Documents operations

Available operations for managing physical inventory documents:

- **Get PhysInventory Doc Item** - Reads information on physical inventory items

Operation ID: getPhysInventoryDocItem

SAP Product Master operations

Available operations for managing product master data:

- **Get Product Master Item** - Returns product master records

Operation ID: getProductMaster

- **Get Plant Data By Material** - Returns plant data of product master record

Operation ID: getPlantDataByMaterial

Required parameter: Product (string, max 40 characters)

- **Get Supply Planning Data By Material** - Returns supply planning data by product number and plant

Operation ID: getSupplyPlanningDataByMaterial

Required parameters: Product (string, max 40 characters), Plant (string, max 4 characters)

Query parameters

SAP connectors support standard query parameters to filter, sort, and format API responses. Use these parameters to optimize data retrieval and processing.

Supported Query Parameters

#	Parameter	Description	Type
1	\$top	Limits the number of returned items	integer
2	\$skip	Skips the specified number of items	integer
3	\$filter	Filters results based on specified criteria	string
4	\$orderby	Orders results by specified fields	array
5	\$select	Selects specific properties to return	array
6	\$expand	Expands related entities	array
7	\$inlinecount	Includes count of items in response	string

Manage SAP workload integrations

After you create your SAP workload integration, you can manage it through several options.

Edit integration settings

Follow these steps to modify your SAP workload integration settings.

1. In the Amazon Quick console, choose **Connectors**.
2. Choose your SAP workload connector from the connector list.
3. Select your integration from the list and choose **Edit**.
4. Modify integration settings:

- Update authentication credentials (username and password).
 - Change SAP system URL or connection settings.
 - Modify integration name or description.
5. Choose **Save changes** to apply your modifications.

Share integration

You can share SAP workload connectors with other users in your organization.

1. From the SAP integration details page, choose **Share**.
2. Configure sharing options:
 - **Share with specific users** - Enter user names or email addresses.
 - **Share with organization** - Make available to all users in your organization.
3. Set permission levels for shared access.
4. Choose **Share integration** to apply the sharing settings.

Delete integration

Follow these steps to permanently remove your SAP integration.

1. From the SAP workload integration details page, choose **Delete**.
2. Review the deletion impact, including any workflows or automations using this integration.
3. Type the integration name to confirm deletion.
4. Choose **Delete integration** to permanently remove it.

Troubleshoot SAP workload integrations

Use these troubleshooting tips to resolve common SAP workload integration issues.

Authentication issues

OAuth 2.0 authentication failures

Symptoms: Token generation fails, invalid client credentials, or OAuth scope errors.

Resolution:

- Verify OAuth client ID and client secret are correct
- Check that OAuth is properly configured in SAP using transaction /IWFND/MAINT_SERVICE
- Ensure the required scopes are properly configured for your SAP workload connector type
- Verify the token URL format matches your SAP system configuration

Basic authentication failures

Symptoms: Login failures, invalid credentials, or access denied errors.

Resolution:

- Verify SAP username and password are correct
- Check that the user account has necessary SAP authorizations
- Ensure the user account is not locked or expired
- Verify the domain URL is accessible and correctly formatted

SAP system configuration issues

API service not activated

Symptoms: Service unavailable errors, API endpoint not found, or HTTP 404 responses.

Resolution:

- Verify the required SAP API service is activated in your system
- Check API service status using SAP transaction codes
- Ensure the API service corresponds to your chosen SAP connector type
- Contact your SAP administrator to activate missing API services

Connection timeouts

Symptoms: Request timeouts, network connectivity errors, or slow response times.

Resolution:

- Check that your SAP system URL is correct and accessible
- Verify network connectivity allows connections to SAP system
- Ensure firewall rules permit the integration traffic

- Check SAP system performance and availability

Permission and authorization errors

Insufficient SAP authorizations

Symptoms: Access denied errors, missing authorization messages, or restricted operation failures.

Resolution:

- Ensure the authenticated user has required SAP authorizations for the specific module
- Verify the user has access to necessary transaction codes
- Check that the user account has appropriate role assignments
- Contact your SAP administrator to grant missing permissions

Scope permission errors

Symptoms: OAuth scope errors, insufficient permissions for API operations, or restricted access messages.

Resolution:

- Verify the OAuth configuration includes the required scope for your connector type
- Check that the scope permissions are properly granted in SAP system
- Ensure the OAuth client has been granted the necessary API access rights

Data format and parameter errors

Invalid parameter formats

Symptoms: Data validation errors, invalid field length messages, or parameter format exceptions.

Resolution:

- Review action parameters and ensure they match expected SAP data formats
- Check field lengths match SAP system requirements (e.g., BusinessPartner max 10 characters)
- Verify data types are correct for the specific SAP module

- Ensure all required fields are provided for the operation

Query parameter errors

Symptoms: Query syntax errors, unsupported parameter messages, or malformed request errors.

Resolution:

- Verify parameters use correct syntax (e.g., `$filter`, `$top`, `$skip`)
- Check that the parameter values are properly formatted
- Ensure the SAP API supports the specific parameters being used
- Refer to SAP API documentation for supported query options

SAP system availability issues

SAP system unavailability

Symptoms: Connection refused errors, system not responding, or service unavailable messages.

Resolution:

- Check SAP system status and availability with your SAP administrator
- Verify if there are scheduled maintenance windows affecting the system
- Check for any SAP system alerts or known issues
- Retry the operation after confirming system availability

ServiceNow integration

Use the ServiceNow integration to perform actions within your ServiceNow instances, including managing incidents, problems, change requests, knowledge base articles, and attachments. This integration uses the ServiceNow REST API. For more information, see [REST API](#) in the ServiceNow documentation.

Setting up this integration involves two steps. First, you configure an OAuth application in your ServiceNow instance. Then, you create the integration in Amazon Quick and connect it to your ServiceNow app. For information about the authentication methods that Amazon Quick supports, see [the section called "Authentication methods"](#).

Before you begin

Before you set up the integration, verify that you have the following.

- A ServiceNow instance. This integration is validated against the Xanadu release.
- A ServiceNow user account with permissions to create OAuth applications (admin role required).
- For service authentication (client credentials), your instance must be running the Washington DC release or later.
- For subscription requirements, see [the section called "Set up integrations in the console"](#).

Configure ServiceNow OAuth

Before you configure Amazon Quick, create an OAuth application endpoint in your ServiceNow instance. Complete all of the following steps in ServiceNow before moving to the Amazon Quick console.

For more information, see [Create an endpoint for clients to access the instance](#) in the ServiceNow documentation.

Register the OAuth application

To register the OAuth application, complete the following steps.

1. In your ServiceNow instance, navigate to **All > System OAuth > Application Registry** and choose **New**.
2. Choose **Create an OAuth API endpoint for external clients**.
3. Complete the form:
 - **Name** – A descriptive name for the OAuth application.
 - **Redirect URL** – `https://{region}.quicksight.aws.amazon.com/sn/oauthcallback`

Replace `{region}` with your AWS Region (for example, `us-east-1`).
4. Choose **Submit**.
5. Reopen the application registry entry and choose the lock icon next to **Client Secret** to reveal the value.
6. Copy the **Client ID** and **Client Secret** values. You need these when you configure the integration in Amazon Quick.

Additional steps for service authentication (client credentials)

If you plan to use service authentication, complete these additional steps after registering the OAuth application. The client credentials grant type was introduced in the ServiceNow Washington DC release. For more information, see [Up Your OAuth2.0 Game: Inbound Client Credentials with Washington DC](#) in the ServiceNow Community.

1. Enable the client credentials grant type. Navigate to `sys_properties.list` using the filter navigator and create a new system property with the following values:
 - **Name** – `glide.oauth.inbound.client.credential.grant_type.enabled`
 - **Type** – `true | false`
 - **Value** – `true`
2. Verify that the following plugins are installed (navigate to **Admin > Application Manager**):
 - OAuth 2.0 (`com.snc.platform.security.oauth`)
 - REST API Provider (`com.glide.rest`)
 - Authentication scope (`com.glide.auth.scope`)
 - REST API Auth Scope Plugin (`com.glide.rest.auth.scope`)
3. Return to your OAuth application in **System OAuth > Application Registry**. Add the **OAuth Application User** field to the form if it isn't visible (use **Configure > Form Builder** to add it).
4. Set the **OAuth Application User** to an appropriately permissioned user, such as a user with the System Administrator role.

Important

With service authentication, all actions execute as the configured OAuth application user. Any Amazon Quick user with access to this integration can perform actions using that account's permissions. Configure the account permissions to match your organization's security requirements.

Set up the integration in Amazon Quick

After you complete the ServiceNow OAuth configuration, create the integration in Amazon Quick.

1. In the Amazon Quick console, choose **Connectors**.

2. Choose the **Create for your team** tab.
3. Find and choose **ServiceNow**.
4. Enter the integration details:
 - **Name** – Descriptive name for your ServiceNow integration.
 - **Description** (Optional) – Purpose of the integration.
5. Choose your connection type and fill in the connection settings:
 - a. For **User authentication (OAuth)**, configure the following fields:
 - **Base URL** – `https://{your-instance}.service-now.com`
 - **Client ID** – Client ID from your ServiceNow OAuth application.
 - **Client Secret** – Client secret from your ServiceNow OAuth application.
 - **Token URL** – `https://{your-instance}.service-now.com/oauth_token.do`
 - **Auth URL** – `https://{your-instance}.service-now.com/oauth_auth.do`
 - **Redirect URL** – `https://{region}.quicksight.aws.amazon.com/sn/oauthcallback`
 - b. For **Service authentication (client credentials)**, configure the following fields:
 - **Authentication type** – Service-to-service OAuth
 - **Base URL** – `https://{your-instance}.service-now.com`
 - **Client ID** – Client ID from your ServiceNow OAuth application.
 - **Client Secret** – Client secret from your ServiceNow OAuth application.
 - **Token URL** – `https://{your-instance}.service-now.com/oauth_token.do`
6. Choose **Create and continue**.
7. Choose users to share the integration with.
8. Choose **Next**.

For user authentication, navigate to **Connectors** > your ServiceNow connector name, and choose **Sign in** to complete the OAuth authorization flow.

Available actions

After you set up the integration, the following actions are available.

ServiceNow available actions

Category	Action	Description
Incidents	List Incidents	Retrieve existing incidents.
Incidents	Create Incident	Create an incident record to document a deviation from an expected standard of operation.
Incidents	View Incident	Retrieve the details of a specific incident.
Incidents	Update Incident	Update an incident record.
Incidents	Delete Incident	Delete an incident.
Problems	List Problems	Retrieve existing problems.
Problems	Create Problem	Create a new problem record.
Problems	View Problem	Retrieve the details of a specific problem record.
Problems	Update Problem	Update a problem record.
Problems	Delete Problem	Delete a problem.
Change requests	List Change Requests	Retrieve all change requests.
Change requests	Create Change Request	Create a change request to implement a controlled process for modifying approved and supported configuration items (CIs).
Change requests	View Change Request	Retrieve detailed information about a specific change request.
Change requests	Update Change Request	Modify a change request.
Change requests	Delete Change Request	Delete a change request.

Category	Action	Description
Knowledge base articles	Create Knowledge Base Article	Create a knowledge base article. Requires the Knowledge API (sn_km_api) plugin.
Knowledge base articles	Update Knowledge Base Article	Modify a knowledge base article.
Knowledge base articles	Delete Knowledge Base Article	Delete a knowledge base article.
Attachments	Retrieve Attachments Metadata	Retrieve metadata for attachment files.
Attachments	Retrieve Attachment Metadata	Retrieve metadata for a specific attachment file.
Attachments	Retrieve Attachment Content	Retrieve the binary file attachment content.
Attachments	Upload Binary Attachment	Upload a binary file as an attachment to a specified record.
Attachments	Upload Multipart Form Attachment	Upload a multipart file attachment.
Attachments	Delete Attachment	Delete an attachment.
Users	List Users	List all user records.
System	List Choices	Retrieve choice list values from the sys_choice table.

 Note

The specific actions available depend on the permissions configured in your ServiceNow instance and the authentication method used.

Limitations

This integration interacts with ServiceNow through the REST API, which does not enforce UI policies, UI actions, or client scripts. These rules only apply in the ServiceNow browser interface. Server-side business rules, ACLs, and data policies are enforced. For more information, see [REST API](#) in the ServiceNow documentation.

Manage and troubleshoot

To edit, share, or delete your integration, see [the section called “Managing existing integrations”](#).

Authentication issues

- **OAuth authorization fails** – Verify that the Client ID and Client Secret match the values in your ServiceNow Application Registry. Confirm the Redirect URL in ServiceNow matches the URL in your Amazon Quick configuration exactly.
- **Service authentication fails** – Verify that the `glide.oauth.inbound.client.credential.grant_type.enabled` system property is set to `true`. Confirm the **OAuth Application User** field is populated on the application registry record.

Common error messages

- **Actions return permission errors** – Verify that the ServiceNow user or OAuth application user has the required roles to access the target tables (for example, `itil` role for incident management).
- **Connection timeout or unreachable instance** – Verify the Base URL uses the correct ServiceNow instance name. Confirm the ServiceNow instance is accessible and not in maintenance mode.

Shopify integration

With the Shopify connector, you can manage products, orders, customers, and store operations directly in Amazon Quick through natural language.

Shopify uses Custom OAuth app authentication. For more information about the authentication methods that Amazon Quick supports, see [the section called “Authentication methods”](#).

Before you begin

Make sure that you have the following before you set up the integration.

- An active Shopify store with admin access.
- A custom app registered in your Shopify admin or in the Shopify Partners portal, with the Admin API access scopes that you need.
- For Amazon Quick subscription requirements, see [the section called “Set up integrations in the console”](#).

Configuring Shopify

For Custom OAuth app authentication, create a custom app in your Shopify store admin or in the [Shopify Partners portal](#) on the Shopify website. Configure the Admin API access scopes that match the operations you want to use, and add the Amazon Quick callback URL `https://{region}.quicksight.aws.amazon.com/sn/oauthcallback` as an allowed redirection URL. Replace `{region}` with your AWS Region (for example, `us-east-1`). For step-by-step instructions, see [Authentication and authorization](#) in the Shopify developer documentation. Record the Client ID (API key) and Client Secret (API secret key) — you need them when you configure Amazon Quick.

Setting up the connector in Amazon Quick

1. In the Amazon Quick console, choose **Connectors**.
2. Choose the **Create for your team** tab.
3. Find and choose **Shopify**.
4. Enter a **Name** for your connector. Optionally, choose **+ Add Description** to add a description.
5. For **Connection type**, choose **Public network**.
6. For **OAuth Configuration**, choose **Custom OAuth app** and configure the following fields:
 - **Base URL** – The Shopify MCP base URL for your store.
 - **Client ID** – The Client ID (API key) from your Shopify custom app.
 - **Public OAuth client** (Optional) – Select this option if your Shopify custom app is configured as a public client (no client secret).
 - **Client secret** – The Client Secret (API secret key) from your Shopify custom app.

- **Token URL** – The Shopify OAuth token endpoint for your store.
 - **Authorization URL** – The Shopify OAuth authorization endpoint for your store.
 - **Redirect URL** – Pre-filled with the Amazon Quick callback URL.
7. Choose **Next**.
 8. A Shopify authorization window opens. Review the requested permissions and choose **Install app** or the equivalent button.
 9. On the **Review** page, review the available actions for the connector. Choose **Next**.
 10. On the **Publish** page, choose who can access the connector. You can enable access for everyone in your organization or search for specific teams or groups.
 11. Choose **Publish**.

Available actions

After you set up the connector, the actions exposed by Shopify are available. To see the current set of actions for your connector, go to the connector's **Available actions** view in the Amazon Quick console.

Managing and troubleshooting

To edit, share, or delete your connector, see [the section called “Managing existing integrations”](#).

Authentication issues

- **Sign-in fails** – Verify that you can sign in to your Shopify store admin directly. Confirm that the allowed redirection URL in your Shopify custom app matches the Amazon Quick callback URL.
- **Invalid client credentials** – Verify that the Client ID (API key) and Client Secret (API secret key) match the values in your Shopify custom app.
- **Insufficient permissions** – Verify that the Admin API access scopes configured on your Shopify custom app match the operations you want to use.

Slack integration

With the Slack connector, you can access the Slack platform directly in Amazon Quick through natural language. You can send messages, manage channels, search content, and interact with users and user groups without leaving Amazon Quick.

Amazon Quick supports multiple authentication methods for Slack. Choose the method that best fits your organization's security requirements.

- **Default OAuth app** – Uses an AWS-managed OAuth application. No additional credentials are needed. You authenticate directly with your Slack workspace.
- **Custom OAuth app** – Uses a customer-managed Slack app. This option gives your organization full control over the OAuth configuration.
- **Bearer Token** – Uses a Slack bot token for authentication. This method is suitable for bot-level access to a workspace.

For more information about the authentication methods that Amazon Quick supports, see [the section called “Authentication methods”](#).

Before you begin

Make sure that you have the following before you set up the integration.

- A Slack workspace with access to the channels and data that you want to use.
- For **Custom OAuth app**: Access to the [Slack API dashboard](#) on the Slack website to create an app.
- For **Bearer Token**: A bot token (starts with xoxb-) from a Slack app. For more information, see [Quickstart app settings](#) on the Slack website.
- For Amazon Quick subscription requirements, see [the section called “Set up integrations in the console”](#).

Configuring Slack

If you are using **Default OAuth app** authentication, skip this section and go to [the section called “Setting up the connector in Amazon Quick”](#).

For Custom OAuth app or Bearer Token authentication, complete the applicable steps in Slack before you configure Amazon Quick.

Create a Slack app for Custom OAuth app or Bearer Token

Create a Slack app to get the client credentials or bot token that you need for Amazon Quick.

1. Go to the [Slack API dashboard](#) on the Slack website and choose **Create New App**.
2. Choose **From scratch**.

3. Enter an **App Name** and select your Slack workspace. Choose **Create App**.
4. In the left sidebar, choose **OAuth & Permissions**.
5. Under **Redirect URLs**, choose **Add New Redirect URL** and enter the Amazon Quick callback URL: `https://{region}.quicksight.aws.amazon.com/sn/oauthcallback`

Choose **Add**, then choose **Save URLs**.

6. Under **Scopes**, add the Bot Token Scopes that your integration requires. For the recommended scopes, see [the section called "Recommended scopes"](#).
7. In the left sidebar, choose **Basic Information**. Under **App Credentials**, record the following values:
 - **Client ID**
 - **Client secret**
8. For **Bearer Token**: Install the app to your workspace from the **OAuth & Permissions** page. After installation, copy the **Bot User OAuth Token** (starts with xoxb-).

Recommended scopes

The following scopes are requested when you connect to Slack. When you create a custom Slack app, add these scopes on the **OAuth & Permissions** page. For Custom OAuth app, add them as **User Token Scopes**. For Bearer Token, add them as **Bot Token Scopes**.

Slack recommended scopes

Scope	Description
<code>calls:read</code>	Reads call information.
<code>calls:write</code>	Creates and manages calls.
<code>channels:history</code>	Reads message history in public channels.
<code>channels:read</code>	Reads information about public channels.
<code>channels:write</code>	Manages public channels.
<code>chat:write</code>	Sends messages to channels and conversations.
<code>dnd:read</code>	Reads Do Not Disturb settings.

Scope	Description
<code>dnd:write</code>	Updates Do Not Disturb settings.
<code>emoji:read</code>	Reads custom emoji.
<code>files:read</code>	Reads files shared in channels.
<code>files:write</code>	Uploads and manages files.
<code>groups:history</code>	Reads message history in private channels.
<code>groups:read</code>	Reads information about private channels.
<code>groups:write</code>	Manages private channels.
<code>im:history</code>	Reads message history in direct messages.
<code>im:read</code>	Reads information about direct messages.
<code>im:write</code>	Starts direct messages.
<code>links:read</code>	Reads link previews.
<code>links:write</code>	Manages link previews.
<code>mpim:history</code>	Reads message history in group direct messages.
<code>mpim:read</code>	Reads information about group direct messages.
<code>mpim:write</code>	Starts group direct messages.
<code>pins:read</code>	Reads pinned messages.
<code>pins:write</code>	Pins and unpins messages.
<code>reactions:read</code>	Reads emoji reactions on messages.
<code>reactions:write</code>	Adds and removes emoji reactions.
<code>reminders:read</code>	Reads reminders.

Scope	Description
<code>reminders:write</code>	Creates reminders.
<code>remote_files:read</code>	Reads remote files.
<code>remote_files:share</code>	Shares remote files.
<code>search:read</code>	Searches messages and files.
<code>stars:read</code>	Reads starred items.
<code>stars:write</code>	Stars and unstars items.
<code>team:read</code>	Reads workspace information.
<code>usergroups:read</code>	Reads user group information.
<code>usergroups:write</code>	Manages user groups.
<code>users:read</code>	Reads user information.
<code>users:read.email</code>	Reads user email addresses.
<code>users:write</code>	Updates user information.
<code>users.profile:read</code>	Reads user profile information.
<code>users.profile:write</code>	Updates user profile information.

Note

Not all scopes are available for both User Token Scopes and Bot Token Scopes. Some scopes, such as `search:read` and `stars:read`, are available only as User Token Scopes. When you configure scopes for Bearer Token authentication, verify scope availability in the [Slack API scopes reference](#) on the Slack website.

Setting up the connector in Amazon Quick

Connect from the Available tab

If you want to use Default OAuth app authentication, you can connect directly from the **Available** tab without additional configuration.

1. In the Amazon Quick console, choose **Connectors**.
2. On the **Available** tab, find **Slack** and choose **Connect**.
3. Complete the Slack sign-in flow and grant the requested permissions.

To configure a connector with Custom OAuth app or Bearer Token instead, use the **Create for your team** tab as described below.

Create from the Create for your team tab

After you complete any required Slack configuration, create the connector in Amazon Quick.

1. In the Amazon Quick console, choose **Connectors**.
2. Choose the **Create for your team** tab.
3. Find and choose **Slack**.

Note

If a Slack connector already exists, a dialog appears with your existing connectors. To use an existing connector, choose it. To create a new one, choose **No, create new**.

4. Enter a **Name** for your connector. Optionally, choose **+ Add Description** to add a description.
5. For **Connection type**, choose **Public network**.
6. For **OAuth Configuration**, choose one of the following authentication methods and configure the required fields.
 - a. For **Default OAuth app**:

No additional credentials are needed. Choose **Next** to continue.
 - b. For **Custom OAuth app**, configure the following fields:
 - **Base URL** (Optional) – The Slack API base URL. Example: `https://slack.com/api`

- **Client ID** – The client ID from your Slack app.
 - **Client secret** – The client secret from your Slack app.
 - **Token URL** – The token endpoint. Example: `https://slack.com/api/oauth.v2.access`
 - **Authorization URL** – The authorization endpoint. Example: `https://slack.com/oauth/v2/authorize`
 - **Redirect URL** – Pre-filled with the Amazon Quick callback URL.
- c. For **Bearer Token**, configure the following fields:
- **Base URL** (Optional) – The Slack API base URL. Example: `https://slack.com/api`
 - **Bot token** – Your Slack bot token (starts with xoxb-).
7. Choose **Next**.
 8. If you chose **Default OAuth app** or **Custom OAuth app**, a Slack authorization window opens. Review the requested permissions, select your workspace, and choose **Allow**.
 9. On the **Review** page, review the available actions for the connector. Choose **Next**.
 10. On the **Publish** page, choose who can access the connector. You can enable access for everyone in your organization or search for specific teams or groups.
 11. Choose **Publish**.

Available actions

After you set up the connector, the following actions are available.

Slack available actions

Category	Action	Description
Messages	Chat Post Message	Posts a message to a channel, direct message, or private group.
Messages	Send Me Message	Sends a third-person action message to a channel.
Messages	Get Message Permalink	Retrieves a permalink URL for a specific message.

Category	Action	Description
Channels	Conversations Create	Creates a public or private channel.
Channels	Conversations History	Retrieves message history from a channel. Returns main timeline messages only, not threaded replies.
Channels	Conversations Replies	Retrieves threaded replies to a specific parent message.
Channels	Conversations Invite	Invites users to an existing channel.
Channels	Conversations Join	Joins an existing channel.
Channels	Conversations Kick	Removes a user from a channel.
Channels	Conversations Leave	Leaves a channel.
Channels	Conversations Members	Lists member IDs for a channel.
Channels	Conversations Open	Opens or resumes a direct message or multi-person direct message.
Channels	Conversations Set Topic	Sets or updates the topic for a channel.
Channels	Get Conversation	Retrieves metadata for a channel by ID.
Channels	List Channels	Lists channels available to the user.
Channels	List Conversations For User	Lists channels and direct messages accessible to a user.
Reactions	Create Reaction	Adds an emoji reaction to a message.
Reactions	Delete Reaction	Removes an emoji reaction from a message.
Reactions	List Reactions	Lists reactions on a message, file, or file comment.

Category	Action	Description
Files	Upload Or Create File	Uploads files or media to channels or threads.
Files	Upload To URL External	Adds a reference to an external file for sharing in Slack.
Files	Download File	Downloads a file shared in Slack.
Files	Files Info	Retrieves metadata and comments for a file.
Files	Files List	Lists files in a workspace with optional filters.
Search	Search All	Searches messages and files across the workspace.
Search	Search Messages	Searches messages with date ranges and filters.
Search	Search Channels	Searches channels by name, topic, or purpose.
Search	Search Users	Searches users by email, name, or display name.
Search	Search Context	Searches across messages, files, channels, and users. Supports semantic search on workspaces with Slack AI.
Search	Get Search Configuration	Checks whether semantic search is available on the workspace.
Users	Get User Details	Retrieves information for a user by ID.
Users	Get User Profile	Retrieves profile information for a user.
Users	Users Profile Set	Updates a user's profile fields.

Category	Action	Description
Users	Users List	Lists all users in the workspace.
User groups	User Groups Create	Creates a new user group.
User groups	User Groups List	Lists user groups in the workspace.
User groups	User Groups Users Update	Replaces all members of a user group.
Reminders	Reminders Add	Creates a reminder with specified text and time.
Workspace	Get Team	Retrieves metadata about the Slack workspace.

Note

The actions that you can use depend on the scopes configured for your Slack app and the channels accessible to the authenticated user.

Managing and troubleshooting

To edit, share, or delete your connector, see [the section called “Managing existing integrations”](#).

Authentication issues

- **Sign-in fails (Default OAuth app or Custom OAuth app)** – Verify that your Slack workspace is active and that you can sign in to [the Slack website](#) directly. For Custom OAuth app, confirm that the redirect URL in your Slack app matches the Amazon Quick callback URL.
- **Invalid client credentials (Custom OAuth app)** – Verify that the Client ID and Client secret match the values in your Slack app's Basic Information section.
- **Bearer Token rejected** – Verify that your bot token is active and starts with xoxb-. Tokens can be revoked from the Slack API dashboard.
- **Insufficient permissions** – Verify that the scopes configured for your Slack app include the permissions required for the actions that you want to use. See [the section called “Recommended scopes”](#).

- **Channel not found or not in channel** – Verify that the bot has been added to the channel. Invite the bot to the channel before attempting to post messages or read history.

Smartsheet integration

With the Smartsheet connector, you can manage sheets, rows, reports, and search across your Smartsheet workspaces directly in Amazon Quick through natural language.

Prerequisites

Before you set up the integration, make sure that you meet the following requirements.

- A Smartsheet account with a Business, Enterprise, or Advanced Work Management plan. Free accounts cannot generate API access tokens or register OAuth apps. For more information, see [Smartsheet pricing](#) on the Smartsheet website and [Smartsheet OAuth](#) in the Smartsheet documentation.
- [Smartsheet Developer Tools](#) activated for your account.
- For subscription requirements, see [the section called “Set up integrations in the console”](#).

Configure Smartsheet Developer Tools

Before you configure Amazon Quick, set up credentials in Smartsheet. Each user signs in with their own Smartsheet account, and actions run with that user's permissions. This method uses Custom OAuth app. Complete the [the section called “Register for Developer Tools and create a developer profile”](#) and [the section called “Register the OAuth application”](#) sections below.

For more information about authentication methods, see [the section called “Authentication methods”](#).

Register for Developer Tools and create a developer profile

1. Go to the [Developer Tools Registration](#) page and register the Smartsheet account you want to use with your apps. For more information, see [Register for Developer Tools](#) in the Smartsheet API documentation.
2. After Smartsheet activates Developer Tools, sign in to the Smartsheet application and choose your **Account** icon in the lower-left corner, then choose **Developer Tools**.
3. Choose **Create Developer Profile** and enter a profile name. For more information, see [Create your developer profile](#) in the Smartsheet API documentation.

 Tip

Smartsheet recommends using a dedicated service account for OAuth apps rather than a personal account.

Register the OAuth application

1. In Smartsheet Developer Tools, choose **Create New App**.
2. Complete the form:
 - **App name** – A name to identify your app to users.
 - **App description** – A brief description of the integration.
 - **App URL** – The URL that launches your app, or a landing page.
 - **App contact/support** – Support contact information.
 - **App redirect URL** – `https://{region}.quicksight.aws.amazon.com/sn/oauthcallback`

Replace `{region}` with your AWS Region (for example, `us-east-1`).
3. Choose **Save**. Smartsheet generates the **App client ID** and **App secret**.
4. Copy the **Client ID** and **Client Secret** values. You need these when you configure the integration in Amazon Quick.

For more information, see [Register an app](#) in the Smartsheet API documentation.

 Note

When you configure the OAuth app, the connector requests access scopes from Smartsheet that determine what it can do on behalf of the authenticated user. Access scopes don't override existing sharing permissions. For more information, see [Access scopes](#) in the Smartsheet API documentation.

Set up the integration in Amazon Quick

1. In the Amazon Quick console, choose **Connectors**.

2. Choose the **Create for your team** tab.
3. Find and choose **Smartsheet**.
4. Enter the integration details:
 - **Name** – Descriptive name for your Smartsheet integration.
 - **Description** (Optional) – Purpose of the integration.
 - **Base URL** – `https://mcp.smartsheet.com`
 - **Client ID** – App client ID from your Smartsheet Developer Tools app registration.
 - **Client Secret** – App secret from your Smartsheet Developer Tools app registration.
 - **Token URL** – `https://api.smartsheet.com/2.0/token`
 - **Auth URL** – `https://app.smartsheet.com/b/authorize`
5. Choose **Create and continue**.
6. Choose users to share the integration with.
7. Choose **Next**.

Note

If you're using a Smartsheet regional server outside of the US, replace the MCP server URL with one for your region:

- For Europe, use `https://mcp.smartsheet.eu`
- For Australia, use `https://mcp.smartsheet.au`

Available actions

After you set up the connector, the actions exposed by the Smartsheet MCP server are available. To see the current set of actions for your connector, open the connector's **Available actions** view in the Amazon Quick console. For a full list of available MCP tools, see [MCP Server Tools](#) in the Smartsheet documentation.

Managing and troubleshooting

To edit, share, or delete your connector, see [the section called “Managing existing integrations”](#).

Authentication issues

- **OAuth authorization fails** – Verify that the Client ID and Client Secret match the values in your Smartsheet Developer Tools app registration. Confirm the redirect URL in Smartsheet matches the URL in your Amazon Quick configuration exactly. For a list of OAuth error types, see [OAuth error types](#) in the Smartsheet API documentation.
- **Developer Tools not available** – Verify that Developer Tools is activated for your Smartsheet account. Free accounts do not support Developer Tools. If your request was denied, contact your Smartsheet Customer Success Manager.

Common error messages

- **Actions return permission errors** – Verify that the authenticated user has the required sharing permissions on the target sheets. OAuth scopes don't override sharing-level access controls. For more information, see [Resource access levels](#) in the Smartsheet API documentation.
- **Sheet not found** – Verify the sheet ID is correct and that the authenticated user has at least viewer access to the sheet.
- **API rate limit errors** – The Smartsheet API enforces rate limits. For more information, see [Limitations](#) in the Smartsheet API documentation.

Snowflake Cortex Agent integration

Note

This guide covers an integration validated with a third-party MCP server. Snowflake manages the setup and availability of the Cortex Agent MCP server.

With the Snowflake Cortex Agent connector, you can query Snowflake data and run AI-powered analytics through Cortex Agent directly in Amazon Quick via Model Context Protocol (MCP) server connectivity.

Snowflake Cortex Agent uses MCP server connectivity. To set up this integration, complete the following two steps. First, enable the Cortex Agent MCP server in your Snowflake account and obtain the server endpoint URL. Then, create the integration in Amazon Quick. For more information about Snowflake's Cortex Agent MCP server, see [Cortex Agents MCP server](#) in the Snowflake documentation.

Prerequisites

Before you set up the integration, make sure that you have the following resources and access.

- An active Snowflake account with Cortex Agent enabled and the Cortex Agent MCP server configured. For information about enabling the Cortex Agent MCP server, see [Cortex Agents MCP server](#) in the Snowflake documentation.
- The MCP server endpoint URL from your Snowflake account.
- Appropriate roles, warehouses, databases, and schemas configured in your Snowflake account for the operations you want to perform.
- For information about subscription requirements, see [the section called “Set up integrations in the console”](#).

Set up the integration in Amazon Quick

After you verify your Snowflake account and obtain the MCP server endpoint, create the integration in Amazon Quick.

1. In the Amazon Quick console, choose **Connectors**.
2. Choose the **Create for your team** tab.
3. Find and choose **Snowflake Cortex Agent**.
4. On the **Connect to Snowflake Cortex Agent** page, enter the following fields:
 - **Name** – A descriptive name for your Snowflake Cortex Agent connector.
 - **Description** (Optional) – Notes about how you plan to use this connection.
 - **MCP server endpoint** – The Cortex Agent MCP server endpoint URL from your Snowflake account. Example: `https://mcp.example.com/sse`.
 - **Connection type** – Choose **Public network**.
5. Choose **Next**.
6. If the Cortex Agent MCP server requires authentication, Amazon Quick uses MCP OAuth discovery to handle the sign-in flow. If a Snowflake sign-in popup appears, enter your Snowflake credentials and grant the requested permissions. For more information about how Amazon Quick handles MCP authentication, see [the section called “Prepare MCP server setup and authentication”](#).

7. On the **Review** page, Amazon Quick discovers the tools that Cortex Agent exposes and displays them in a table. Review the list of available actions and choose **Next**.
8. (Optional) On the **Share integration** page, search for teams or groups to share the integration with.
9. Choose **Done**.

The integration appears in the **Existing actions** panel with a status of **Available**.

Available actions

After you set up the integration, the actions exposed by Snowflake Cortex Agent are available. To see the current set of actions for your connector, go to the connector's **Available actions** view in the Amazon Quick console.

Manage and troubleshoot

To edit, share, or delete your integration, see [the section called "Managing existing integrations"](#). For general MCP integration troubleshooting, see [the section called "Troubleshooting"](#).

Visier Agent integration

With the Visier Agent connector, you can access Visier's people analytics platform directly in Amazon Quick through natural language. You can query live workforce data, analyze headcount trends, assess retention health, generate workforce briefings, and more without leaving Amazon Quick.

Visier Agent uses Custom OAuth app authentication. For more information about the authentication methods that Amazon Quick supports, see [the section called "Authentication methods"](#).

Before you begin

Make sure that you have the following before you set up the integration.

- An active Visier account with API access and Visier Agent enabled. For information about enabling Visier Agent and configuring authentication, see [Visier MCP Server documentation](#).
- Appropriate data access scopes configured in your Visier account for the workforce data you want to query.

- OAuth credentials (Client ID and Client Secret) from your Visier administrator.
- For Amazon Quick subscription requirements, see [the section called “Set up integrations in the console”](#).

Configuring Visier

Before you configure Amazon Quick, complete the following steps in your Visier admin console.

1. Sign in to your Visier admin console and confirm that your account is active.
2. Enable the Visier Agent capability under **Settings, API & Integrations**.
3. Register an OAuth client for Amazon Quick and add the Amazon Quick callback URL to the list of redirect URLs: `https://{region}.quicksight.aws.amazon.com/sn/oauthcallback`

Replace `{region}` with your AWS Region (for example, `us-east-1`).

4. Record the **Client ID** and **Client Secret**. You need them when you configure Amazon Quick.
5. Confirm that the workforce data and analytics you want to query are accessible to your account.

Setting up the connector in Amazon Quick

1. In the Amazon Quick console, choose **Connectors**.
2. Choose the **Create for your team** tab.
3. Find and choose **Visier Agent**.
4. Enter a **Name** for your connector. Optionally, choose **+ Add Description** to add a description.
5. For **Connection type**, choose **Public network**.
6. For **OAuth Configuration**, choose **Custom OAuth app** and configure the following fields:
 - **Base URL** – The Visier Agent MCP base URL, pre-filled by Amazon Quick.
 - **Client ID** – The OAuth client ID from your Visier admin console.
 - **Public OAuth client** (Optional) – Select this option if your Visier OAuth client is configured as a public client (no client secret).
 - **Client secret** – The OAuth client secret from your Visier admin console.
 - **Token URL** – The OAuth token endpoint, pre-filled by Amazon Quick.
 - **Authorization URL** – The OAuth authorization endpoint, pre-filled by Amazon Quick.

- **Redirect URL** – Pre-filled with the Amazon Quick callback URL.
7. Choose **Next**.
 8. In the Visier sign-in popup that appears, enter your Visier credentials and choose **Sign in** to authorize Amazon Quick to access Visier on your behalf.
 9. On the **Review** page, Amazon Quick discovers the tools that Visier Agent exposes and displays them in a table. Review the list of available actions and choose **Next**.
 10. On the **Publish** page, choose who can access the connector. You can enable access for everyone in your organization or search for specific teams or groups.
 11. Choose **Publish**.

Available actions

After you set up the connector, you can use the following actions. Visier Agent exposes its people analytics capabilities as callable tools.

Visier Agent available actions

Action	Description
ask_vee_question	Asks Visier's Vee AI a natural language workforce question, such as headcount, turnover, or requisitions.
aggregate_metric_values	Computes aggregate values for a metric, optionally grouped by a dimension and filtered by specific members.
sample_vee_questions	Returns a list of sample questions that you can ask Vee to get started.
search_analytic_objects	Finds analytic objects (for example, Employee or Applicant) matching a search string.
search_analytic_object_properties	Returns all available properties for a given analytic object.
list_analytic_object_property_values	Lists property values for specific analytic object members, such as employee tenure or pay.
search_metrics	Finds metrics matching a search string.

Action	Description
search_dimensions_for_metric	Finds dimensions available for a given metric, usable for grouping or filtering.
search_dimensions_for_analytic_object	Finds dimensions available for a given analytic object.
search_filter_dimension_members	Returns filterable members for a given dimension.
get_object_reference_graph	Returns the reference graph showing how a given analytic object connects to others.

Managing and troubleshooting

To edit, share, or delete your connector, see [the section called “Managing existing integrations”](#).

Authentication issues

- **MCP server connection fails** – Verify that you enabled the Visier Agent capability in your Visier admin console under **Settings, API & Integrations**. Confirm that you entered the Base URL, Token URL, and Authorization URL correctly in Amazon Quick.
- **Invalid client credentials** – Verify that the Client ID and Client secret match the values from your Visier admin console. Confirm that the redirect URI in your Visier OAuth client matches the Amazon Quick callback URL.
- **Data access errors** – Verify that the data access scopes configured in your Visier account include the workforce data you are trying to query. Contact your Visier administrator to review your account permissions.
- **No tools discovered** – If Amazon Quick does not discover any tools after you create the connector, confirm that Visier Agent is active and that your authentication credentials have not expired.

Web Crawler integration

Applies to: Enterprise Edition

With Web Crawler integration in Amazon Quick, you can create knowledge bases from website content by crawling and indexing web pages. This integration supports data ingestion capabilities with different authentication options.

Web Crawler capabilities

You can ask questions about content stored on websites and web pages. For example, you can search documentation sites, knowledge bases, or specific information across multiple web pages.

With this integration, you can access and understand web content regardless of location or type. You also get contextual details such as publication dates, modification history, and page ownership for more efficient information discovery.

Note

Web Crawler integration supports data ingestion only. It doesn't provide action capabilities for managing websites or web services.

Prerequisites

Before you set up Web Crawler integration, make sure you have the following:

- Website URLs to crawl and index.
- You need a Amazon Quick account with Enterprise Edition.
- A website that is not behind a firewall and does not require special browser plugins to connect.

Prepare website access and authentication

Before setting up the integration in Amazon Quick, prepare your website access credentials. Web Crawler integration supports different authentication methods:

No authentication

Use for crawling websites that don't require authentication.

Basic authentication

Standard HTTP Basic Authentication for secured websites. When you visit a protected site, your browser displays a dialog box that asks for your credentials.

Required credentials:

- **Login page URL** - The URL of the login page
- **Username** - Basic auth username
- **Password** - Basic auth password

Form authentication

For websites that use HTML form-based login pages. You specify XPath expressions to identify the form fields on the login page.

XPath (XML Path Language) is a query language for navigating elements in an HTML or XML document. To find an XPath for a web page element, right-click the element in your browser and choose **Inspect**. In the developer tools, right-click the highlighted HTML code, choose **Copy**, and then choose **Copy XPath**.

Required information:

- **Login page URL** - URL of the login form (for example, `https://example.com/login`)
- **Username** - Login username
- **Password** - Login password
- **Username field XPath** - XPath to username input field (for example, `//input[@id='username']`)
- (Optional) **Username button XPath** - XPath to username button field (for example, `//input[@id='username_button']`)
- **Password field XPath** - XPath to password input field (for example, `//input[@id='password']`)
- **Password button XPath** - XPath to password button (for example, `//button[@type='password']`)

SAML authentication

For websites that use SAML-based single sign-on (SSO) authentication.

SAML (Security Assertion Markup Language) authentication is a federated identity standard that enables SSO. You authenticate through a centralized identity provider (such as Microsoft Azure AD or Okta) instead of entering credentials directly into each application. The identity provider passes a secure token back to the application to grant access.

Required information:

- **Login page URL** - URL of the SAML login page
- **Username** - SAML username
- **Password** - SAML password
- **Username field XPath** - XPath to username input field (for example, `//input[@id='username']`)
- (Optional) **Username button XPath** – XPath to username button field (for example, `//input[@id='username_button']`)
- **Password field XPath** - XPath to password input field (for example, `//input[@id='password']`)
- **Password button XPath** - XPath to password button (for example, `//button[@type='password']`)

XPath configuration examples

Use these XPath examples to configure form and SAML authentication:

Username field examples:

```
//input[@id='username']  
//input[@name='user']  
//input[@class='username-field']
```

Password field examples:

```
//input[@id='password']  
//input[@name='pass']  
//input[@type='password']
```

Submit button examples:

```
//button[@type='submit']  
//input[@type='submit']  
//button[contains(text(), 'Login')]
```

Set up Web Crawler integration

After preparing your website access requirements, create the Web Crawler integration in Amazon Quick.

1. In the Amazon Quick console, choose **Knowledge**.

2. Find **Web Crawler** and choose the **Add (+)** icon.
 3. Choose **Access data from Web Crawler**. Web Crawler integration supports data access only - action execution is not available for web crawling.
 4. Configure integration details and authentication method, then create knowledge bases as needed.
 1. Choose the authentication type for your web crawler integration.
 2. Enter the required details based on your chosen authentication method.
 3. (Optional) Choose a VPC connection to crawl sites hosted in your private network. The VPC connection must be configured in admin settings before you can choose it here. For more information, see [the section called "Setting up a VPC to use with Amazon Quick"](#).
-  **Note**

You can't change the VPC connection after the integration is created. To use a different VPC connection, create a new integration.
4. Choose **Create and continue**.
 5. Enter the name and description for your knowledge base.
 6. Add the content URLs that you want to crawl.
 7. Choose **Create**.

After you choose **Create**, the data sync starts automatically.

Configure crawling

You can configure which websites and pages to crawl and how to filter the content.

Configure URLs and content sources

Configure which websites and pages to crawl:

Direct URLs

Specify individual URLs to crawl:

`https://example.com/docs`

```
https://example.com/blog  
https://example.com/support
```

Limit: Maximum 10 URLs per dataset

Content filters and crawl settings

Crawl scope settings

To view these settings, you must first set up a knowledge base and then examine the advanced settings option.

Crawl depth

- Range: 0-10 (default: 1)
- 0 = crawl only specified URLs
- 1 = include linked pages one level deep
- Higher values follow links deeper into the site

Maximum links per page

- Default: 100
- Maximum: 1,000
- Controls how many links to follow from each page

Wait time

- Default: 1
- The time (in seconds) that the web crawler waits for each page after the page reaches a ready state. Increase this value for pages with dynamic JavaScript content that loads after the main template.

Manage knowledge bases

After setting up your Web Crawler integration, you can create and manage knowledge bases from your crawled website content.

Edit existing knowledge bases

You can modify your existing Web Crawler knowledge bases:

1. In the Amazon Quick console, choose **Knowledge bases**.
2. Choose your Web Crawler knowledge base from the list.
3. Choose the three-dot icon under **Actions**, then choose **Edit knowledge base**.
4. Update your configuration settings as needed and choose **Save**.

Attachments and file crawling

Control whether the system processes files and attachments linked from web pages:

- **Enable file attachment crawling** – Choose this option to crawl and index files and attachments found on web pages, such as PDFs, documents, and media files.

Crawling behavior and sync configuration

Your Web Crawler integration follows these crawling practices:

- **Incremental sync model:** First sync performs full crawl. Subsequent syncs capture changes only.
- **Automatic retry:** Built-in retry logic for failed requests.
- **Duplicate handling:** Automatic detection and deduplication of URLs.
- **Crawler identification:** Identifies itself with the user-agent string `amazon-Quick-on-behalf-of-<UUID>` in request headers. You can target this user-agent in your `robots.txt` file to allow or disallow the crawler specifically. For more information, see [the section called “Robots.txt compliance”](#).

Sitemap discovery

Web Crawler automatically checks for sitemaps by appending common sitemap paths to your seed URLs. You don't need to provide sitemap URLs separately. The following paths are checked:

```
sitemap.xml
sitemap_index.xml
sitemap/sitemap.xml
sitemap/sitemap_index.xml
sitemaps/sitemap.xml
sitemap/index.xml
```

For example, if your seed URL is `https://example.com/docs`, the crawler checks for `https://example.com/docs/sitemap.xml`, `https://example.com/docs/sitemap_index.xml`, and so on.

Note

Web Crawler does not follow recursive sitemap index references. Only the URLs listed directly in a discovered sitemap are used. Sitemap directives in `robots.txt` are not used for sitemap discovery.

Robots.txt compliance

Web Crawler respects the `robots.txt` protocol and honors user-agent and allow or disallow directives. Use these directives to control how the crawler accesses your site.

How robots.txt checking works

- **Host-level checking:** Web Crawler reads `robots.txt` files at the host level (for example, `example.com/robots.txt`).
- **Multiple host support:** For domains with multiple hosts, Web Crawler honors robots rules for each host separately.
- **Fallback behavior:** If Web Crawler can't fetch `robots.txt` due to blocking, parsing errors, or timeouts, it behaves as if `robots.txt` doesn't exist. In this case, the crawler proceeds to crawl the site.

Supported robots.txt fields

Web Crawler recognizes these `robots.txt` fields (field names are case-insensitive, values are case-sensitive):

`user-agent`

Identifies which crawler the rules apply to.

`allow`

A URL path that may be crawled.

disallow

A URL path that may not be crawled.

crawl-delay

The time (in seconds) to wait between requests to your website.

Meta tag support

Web Crawler supports page-level robots meta tags that you can use to control how your data is used. You can specify page-level settings by including a meta tag on HTML pages or in an HTTP header.

Supported meta tags

noindex

Do not index the page. If you don't specify this rule, the page may be indexed and eligible to appear in experiences.

nofollow

Do not follow the links on this page. If you don't specify this rule, Web Crawler may use the links on the page to discover those linked pages.

You can combine multiple values using a comma (for example, "noindex, nofollow").

Note

To detect meta tags, Web Crawler must access your page. Don't block your page with robots.txt, because this prevents the page from being recrawled.

Troubleshooting

Use this section to resolve common issues with Web Crawler integration.

Authentication failures

Symptoms:

- "Unable to authenticate" error messages

- 401/403 HTTP responses
- Login page redirect loops
- Session timeout errors

Resolution steps:

1. Verify the site is reachable from the AWS Region where the Amazon Quick instance is set up.
2. Verify that your credentials are correct and haven't expired.
3. Check authentication endpoint availability and accessibility.
4. Validate XPath configurations by testing them in browser developer tools.
5. Review browser network logs to understand the authentication flow.
6. Ensure the login page URL is correct and accessible.
7. Test authentication manually using the same credentials.

Access and connectivity issues**Symptoms:**

- Connection timeouts and network errors
- Network unreachable errors
- DNS resolution failures

Resolution steps:

1. Verify network connectivity to target websites.
2. Validate site accessibility:
 - Check DNS resolution for target domains.
 - Verify SSL/TLS configuration and certificates.
 - Test access from different networks if possible.

DNS resolution

The Web Crawler uses DNS to resolve website hostnames (for example, `www.example.com`) to IP addresses. By default, it uses public DNS resolution.

When crawling sites inside a VPC, you may need to configure a private DNS server so the crawler can resolve hostnames for internal sites. Choose one of the following options based on your VPC configuration:

1. **Use the VPC-provided DNS server** — If your VPC has both **DNS hostnames** and **DNS resolution** enabled, you can use the default VPC DNS resolver (typically 10.0.0.2, or more generally the VPC CIDR base+2). For more information, see [the section called "VPC"](#).
2. **Use a custom DNS server** — If your VPC uses a custom DNS resolver, provide the IP address of your organization's internal DNS server. Work with your network administrator to obtain this address.

If you don't configure a DNS server, the crawler resolves only publicly registered hostnames.

JavaScript-dependent navigation

Symptoms:

- Only the seed URL is indexed, no additional pages discovered
- Crawl completes successfully but returns only one document

Resolution steps:

1. Web Crawler executes JavaScript and renders page content, but does not simulate user interactions such as clicks, scrolls, or hover actions. If your site loads navigation links through user interaction (for example, click handlers, infinite scroll, or dynamic menus), the crawler cannot discover those links.
2. Inspect your page in browser developer tools to check if navigation links use standard `<a href="...">` elements. If links are wired through JavaScript event handlers instead, the crawler will not follow them.
3. If your site provides a sitemap, Web Crawler automatically checks for common sitemap paths on your seed URLs. Ensure your sitemap is available at a standard location (for example, `/sitemap.xml`) so the crawler can discover additional URLs without relying on in-page link extraction.
4. Alternatively, provide all target page URLs directly as seed URLs.
5. If content can be exported as HTML, PDF, or text files, consider using the Amazon S3 connector as your data source instead.

Crawl and content issues

Symptoms:

- Missing or incomplete content
- Incomplete crawls or early termination
- Rate limiting errors (429 responses)
- Content not being indexed properly

Resolution steps:

1. Review robots.txt restrictions:
 - Check robots.txt file for crawl restrictions.
 - Verify that the crawler is allowed to access target paths.
 - Ensure robots.txt compliance isn't blocking content.
2. Check rate limiting and throttling:
 - Monitor response headers for rate limit information.
 - Implement appropriate crawl delays.
3. Verify URL patterns and filters:
 - Test regex patterns for accuracy.
 - Check URL formatting and structure.
 - Validate include/exclude pattern logic.
4. Review content restrictions:
 - Check for noindex meta tags on pages.
 - Verify content type support.
 - Ensure content size is within limits.
5. Update the wait time so that the content loads on the page before the crawler starts crawling.

Known limitations

Web Crawler integration has the following limitations:

- **URL limits:** Maximum of 10 seed URLs per dataset. You can't provide sitemap URLs in the seed URL field.
- **Crawl depth:** Maximum crawl depth of 10 levels
- **Security requirements:** HTTPS required for web proxy configurations

The following limitations apply when using the Web Crawler with a VPC connection:

- **No HTTP/3 (QUIC) support:** HTTP/3 is not supported. Most sites will fall back to HTTP/2 automatically, but sites configured for HTTP/3 only will not be accessible.
- **DNS over TCP required:** DNS resolution must use TCP. Verify that your DNS server supports DNS over TCP before configuring VPC crawling.
- **Publicly trusted SSL certificates required:** Internal sites must use a certificate from a well-known certificate authority (for example, Let's Encrypt or DigiCert). Sites using self-signed or private CA certificates fail to connect.
- **IPv4 only:** Only IPv4 addresses are supported. Sites accessible exclusively over IPv6 cannot be crawled.

WhatsApp integration

With the WhatsApp connector, you can send and manage messages through WhatsApp Business directly in Amazon Quick through natural language.

WhatsApp uses Custom OAuth app authentication. For more information about the authentication methods that Amazon Quick supports, see [the section called "Authentication methods"](#).

Before you begin

Make sure that you have the following before you set up the integration.

- An active WhatsApp Business account with access to the WhatsApp Business Platform.
- A registered phone number for sending messages.
- An app registered in the Meta for Developers portal with the WhatsApp Business product configured.
- For Amazon Quick subscription requirements, see [the section called "Set up integrations in the console"](#).

Configuring Meta for Developers

For Custom OAuth app authentication, register an app in [Meta for Developers](#) on the Meta website and add the WhatsApp Business product. Configure your phone number and message templates, and add the Amazon Quick callback URL `https://{region}.quicksight.aws.amazon.com/sn/oauthcallback` as a valid OAuth redirect URI. Replace `{region}` with your AWS Region (for example, `us-east-1`). For step-by-step instructions, see [WhatsApp Cloud API](#) in the Meta for Developers documentation. Record the App ID (Client ID) and App Secret (Client Secret) — you need them when you configure Amazon Quick.

Setting up the connector in Amazon Quick

1. In the Amazon Quick console, choose **Connectors**.
2. Choose the **Create for your team** tab.
3. Find and choose **WhatsApp**.
4. Enter a **Name** for your connector. Optionally, choose **+ Add Description** to add a description.
5. For **Connection type**, choose **Public network**.
6. For **OAuth Configuration**, choose **Custom OAuth app** and configure the following fields:
 - **Base URL** – The WhatsApp MCP base URL.
 - **Client ID** – The App ID from your Meta for Developers app.
 - **Public OAuth client** (Optional) – Select this option if your Meta app is configured as a public client (no client secret).
 - **Client secret** – The App Secret from your Meta for Developers app.
 - **Token URL** – The Meta OAuth token endpoint.
 - **Authorization URL** – The Meta OAuth authorization endpoint.
 - **Redirect URL** – Pre-filled with the Amazon Quick callback URL.
7. Choose **Next**.
8. A Meta authorization window opens. Review the requested permissions and choose **Continue**.
9. On the **Review** page, review the available actions for the connector. Choose **Next**.
10. On the **Publish** page, choose who can access the connector. You can enable access for everyone in your organization or search for specific teams or groups.
11. Choose **Publish**.

Available actions

After you set up the connector, the actions exposed by WhatsApp are available. To see the current set of actions for your connector, go to the connector's **Available actions** view in the Amazon Quick console.

Managing and troubleshooting

To edit, share, or delete your connector, see [the section called “Managing existing integrations”](#).

Authentication issues

- **Sign-in fails** – Verify that your Meta for Developers app is active and that the WhatsApp Business product is configured. Confirm that the OAuth redirect URI in your Meta app matches the Amazon Quick callback URL.
- **Invalid client credentials** – Verify that the App ID (Client ID) and App Secret (Client Secret) match the values in your Meta for Developers app.
- **Phone number or template errors** – Verify that the phone number is registered and that any message templates you reference are approved in the WhatsApp Business Platform.

Zapier integration

With the Zapier connector, you can trigger workflows across thousands of applications that Zapier connects to, directly in Amazon Quick through natural language. Use this connector to extend the action capabilities of Amazon Quick to any application that you have already configured in your Zapier account.

Amazon Quick supports multiple authentication methods for Zapier. Choose the method that best fits your organization's security requirements.

- **Default OAuth app** – Uses an AWS-managed OAuth application. No additional credentials are needed. Users authenticate directly with their Zapier account.
- **Custom OAuth app** – Uses a customer-managed OAuth application configured in your Zapier account. This option gives your organization full control over the OAuth configuration.
- **Service-to-Service OAuth** – Uses client credentials for server-to-server authentication without user interaction. Suitable for automated workflows.
- **API Key** – Uses an API key for authentication.

For more information about the authentication methods that Amazon Quick supports, see [the section called "Authentication methods"](#).

Before you begin

Make sure that you have the following before you set up the integration.

- An active Zapier account with the actions or Zaps that you want to expose configured.
- The Amazon Quick integration requires Zapier MCP **Agentic Mode** to work correctly. Managed Mode is not supported.
- For **Custom OAuth app** or **Service-to-Service OAuth**: OAuth credentials from your Zapier MCP configuration.
- For **API Key**: A Zapier-issued API key.
- For Amazon Quick subscription requirements, see [the section called "Set up integrations in the console"](#).

Configuring Zapier

If you are using **Default OAuth app** authentication, skip this section and proceed to [the section called "Setting up the connector in Amazon Quick"](#).

For Custom OAuth app or Service-to-Service OAuth authentication, configure your [Zapier MCP](#) settings on the Zapier website and add the Amazon Quick callback URL `https://{region}.quicksight.aws.amazon.com/sn/oauthcallback` as a redirect URI. Replace `{region}` with your AWS Region (for example, `us-east-1`). Record the Client ID and Client Secret — you need them when you configure Amazon Quick. For API Key authentication, generate an API key from your Zapier account.

Setting up the connector in Amazon Quick

Connect from the Available tab

If you want to use Default OAuth app authentication, you can connect directly from the **Available** tab without additional configuration.

1. In the Amazon Quick console, choose **Connectors**.
2. On the **Available** tab, find **Zapier** and choose **Connect**.
3. Complete the Zapier sign-in flow and grant the requested permissions.

To configure a connector with one of the other authentication methods, use the **Create for your team** tab as described below.

Create from the Create for your team tab

1. In the Amazon Quick console, choose **Connectors**.
2. Choose the **Create for your team** tab.
3. Find and choose **Zapier**.
4. Enter a **Name** for your connector. Optionally, choose **+ Add Description** to add a description.
5. For **Connection type**, choose **Public network**.
6. For **OAuth Configuration**, choose one of the following authentication methods and configure the required fields.

a. For **Default OAuth app**:

No additional credentials are needed. Choose **Next** to continue.

b. For **Custom OAuth app**, configure the following fields:

- **Client ID** – The client ID from your Zapier OAuth app.
- **Public OAuth client** (Optional) – Select this option if your Zapier OAuth app is configured as a public client (no client secret).
- **Client secret** – The client secret from your Zapier OAuth app.
- **Token URL** – The token endpoint. Default: `https://mcp.zapier.com/token`
- **Authorization URL** – The authorization endpoint. Default: `https://mcp.zapier.com/authorize`
- **Redirect URL** – Pre-filled with the Amazon Quick callback URL.

c. For **Service-to-Service OAuth**, configure the following fields:

- **Client ID** – The client ID from your Zapier OAuth app.
- **Client secret** – The client secret from your Zapier OAuth app.
- **Token URL** – The token endpoint. Default: `https://mcp.zapier.com/token`

d. For **API Key**, configure the following fields:

- **API Key** – The Zapier API key.
- **Email** (Optional) – The email address associated with the API key.

7. Choose **Next**.
8. If you chose **Default OAuth app** or **Custom OAuth app**, a Zapier authorization window opens. Review the requested permissions and choose **Allow**.
9. On the **Review** page, review the available actions for the connector. The available actions depend on the Zaps that you have configured in your Zapier account. Choose **Next**.
10. On the **Publish** page, choose who can access the connector. You can enable access for everyone in your organization or search for specific teams or groups.
11. Choose **Publish**.

Available actions

After you set up the connector, the actions that you exposed in your Zapier account are available. Each Zap or action that you authorized in Zapier registers as a callable tool in Amazon Quick. To see the current set of actions for your connector, go to the connector's **Available actions** view in the Amazon Quick console.

Managing and troubleshooting

To edit, share, or delete your connector, see [the section called “Managing existing integrations”](#).

Authentication issues

- **Sign-in fails (Default OAuth app or Custom OAuth app)** – Verify that your Zapier account is active and that you can sign in to Zapier directly. For Custom OAuth app, confirm that the redirect URI in your Zapier OAuth app matches the Amazon Quick callback URL.
- **Invalid client credentials (Custom OAuth app or Service-to-Service OAuth)** – Verify that the Client ID and Client secret match the values in your Zapier OAuth app.
- **API Key authentication fails** – Verify that the API key has not been revoked.
- **Expected actions are missing** – Verify that the Zaps or actions that you want to use are configured and enabled in your Zapier account.

Zendesk Suite integration

Manage your Zendesk support tickets directly from Amazon Quick using natural language. With the Zendesk Suite connector, you can create, update, search, and list support tickets without leaving Amazon Quick.

Zendesk Suite is available as a built-in connector in Amazon Quick. To set up this integration, complete the following two steps. First, prepare your Zendesk account with the required access. Then, create the connector in Amazon Quick and authenticate with your Zendesk credentials.

This connector supports OAuth user authentication and API key service authentication. For information about the authentication methods that Amazon Quick supports, see [the section called "Authentication methods"](#).

Prerequisites

Before you set up the integration, verify that you have the following resources and access.

- An active Zendesk account with an admin role.
- Permissions to access tickets and support data in your Zendesk instance.
- A subscription that meets the requirements described in [the section called "Set up integrations in the console"](#).

Verify your Zendesk account

Before you configure Amazon Quick, verify that your Zendesk account meets the following requirements.

- Your Zendesk account is active and you can sign in to the [Zendesk sign-in page](#) on the Zendesk website.
- The tickets and support data that you want to access are accessible to your account.

Configure Zendesk account

The Custom OAuth app authentication method in Amazon Quick requires an OAuth client configured in Zendesk. The API key authentication method requires an API token.

To configure an OAuth client

1. In your Zendesk instance, navigate to **Admin Center**, then **Apps and integrations**, and then **OAuth clients** under **API**.
2. Choose **Add OAuth Client**.
3. Complete the form with the following fields:

- **Name** – Enter a name for your app.
 - **Description** – Enter a short description.
 - **Identifier** – Enter a unique identifier for the OAuth client.
 - **Redirect URLs** – Enter `https://region.quicksight.aws.amazon.com/sn/oauthcallback`. Replace *region* with your AWS Region, for example, `us-east-1`.
4. Choose **Save**. Zendesk generates a client secret.

 Important

Record the client secret now. You cannot retrieve it after you leave this page.

5. After you capture the secret, choose **Save** again to create the OAuth client.

The Identifier and the secret serve as the Client ID and Client Secret when you configure the connector in Amazon Quick with the Custom OAuth app authentication method.

For more information, see [Using OAuth authentication with your application](#) on the Zendesk website.

To configure an API token

1. Navigate to **Admin Center**, then **Apps and integrations**, and then **API tokens**.
2. Choose **Add API token**.
3. For **Description**, enter a short description for the API token.
4. Choose **Save**. Zendesk generates an API token.

 Important

Record the API token now. You cannot retrieve it after you leave this page.

If you choose the API key authentication method in Amazon Quick, use this API token as the API key value.

For more information, see [Managing API token access to the Zendesk API](#) on the Zendesk website.

Set up the integration in Amazon Quick

After you verify your Zendesk account access, create the connector in Amazon Quick.

1. In the Amazon Quick console, choose **Connectors**.
2. Choose the **Create for your team** tab.
3. Choose **Zendesk Suite** from the options available.
4. In the **Zendesk Suite** dialog, enter the following fields:
 - **Name** – A descriptive name for your Zendesk connector.
 - **Description (Optional)** – Choose **+ Add Description** to add notes about how this connector will be used.
 - **Connection type** – Choose **Public network**.
5. Choose your authentication method and configure the connector settings:
 - **Custom OAuth app** – Enter the following fields for your Zendesk instance:
 - **Base URL**
 - **Client ID**
 - **Client Secret**
 - **Token URL**
 - **Authorization URL**
 - **Redirect URL**
 - **API Key** – Enter the following fields:
 - **Base URL**
 - **API Key** (Zendesk API token)
 - **Email** associated with your Zendesk account
6. Choose **Next**.
7. Review the actions for Zendesk Suite and choose **Next**.
8. Specify who can access the connector:
 - To share with a team, enter a team name or group email.
 - To keep it private, leave the field blank.
 - To share with all users, toggle **Everyone in your organization**.
9. Choose **Publish**.

The connector appears in the **Available** panel with a status of **Connected**.

Available actions

After you set up the connector, you can use the following ticket management actions.

Zendesk Suite available actions

Action	Description
Create Ticket	Create new ticket
List Search Results	Search tickets
List Tickets	View all tickets
Show Ticket	View ticket details
Update Ticket	Edit ticket details

Note

The actions that you can use depend on the permissions that are configured in your Zendesk account.

Troubleshooting

To edit, share, or delete your connector, see [the section called “Managing existing integrations”](#).

Authentication issues

- **Sign-in fails** – Verify that your Zendesk account is active and that you can sign in to the [Zendesk sign-in page](#) on the Zendesk website. If your organization uses single sign-on (SSO), confirm that your identity provider is configured correctly.
- **Access denied** – Confirm that your Zendesk user has an admin role with permissions to access tickets and support data. For API key connections, verify that the API token is valid and the associated email address is correct.

Zoom integration

With the Zoom connector, you can access the Zoom platform directly in Amazon Quick through natural language. You can create and manage meetings, access recordings and meeting summaries, view webinar details, and search contacts without leaving Amazon Quick.

Amazon Quick supports multiple authentication methods for Zoom. Choose the method that best fits your organization's security requirements.

- **Default OAuth app** – Uses an AWS-managed OAuth application. No additional credentials are needed. Users authenticate directly with their Zoom account.
- **Custom OAuth app** – Uses a customer-managed OAuth application registered in the Zoom App Marketplace. This option gives your organization full control over the OAuth configuration.

For more information about the authentication methods that Amazon Quick supports, see [the section called “Authentication methods”](#).

Before you begin

Make sure that you have the following before you set up the integration.

- An active Zoom account with access to the meetings, recordings, and webinars that you want to use.
- For **Custom OAuth app**: Access to the [Zoom App Marketplace](#) on the Zoom website to create an app.
- For Amazon Quick subscription requirements, see [the section called “Set up integrations in the console”](#).

Configuring Zoom

If you are using **Default OAuth app** authentication, skip this section and proceed to [the section called “Setting up the connector in Amazon Quick”](#).

For Custom OAuth app authentication, complete the following steps in the Zoom App Marketplace before configuring Amazon Quick.

Create a Zoom app (Custom OAuth app)

Create an app in the Zoom App Marketplace to obtain the client credentials that you need for Amazon Quick. For more information, see [OAuth](#) on the Zoom Developer website.

1. Sign in to the [Zoom App Marketplace](#) on the Zoom website and go to the app creation portal.
2. Create a new application. Choose **General** for user-based OAuth or **Server-to-Server OAuth** for service-to-service authentication.
3. On the **Basic Information** tab, under **App Credentials**, record the following values. You need them when you configure Amazon Quick.
 - **Client ID**
 - **Client secret**
4. In the **OAuth Information** section, for **Authorized Redirect URI**, add the Amazon Quick callback URL: `https://{region}.quicksight.aws.amazon.com/sn/oauthcallback`
5. Choose the **Scopes** tab and choose **+ Add Scopes** to select the permissions that your integration requires. For the recommended scopes, see [the section called "Recommended scopes"](#).

Recommended scopes

The following scopes are requested when you connect to Zoom. When you create a custom Zoom app, add these scopes on the **Scopes** tab.

Zoom recommended scopes

Scope	Description
meeting:read:list_meetings	Lists meetings.
meeting:read:meeting	Reads meeting details.
meeting:read:summary	Reads meeting summaries.
meeting:read:list_past_participants	Lists past meeting participants.
meeting:read:list_past_instances	Lists past meeting instances.

Scope	Description
meeting:write:meeting	Creates meetings.
meeting:update:meeting	Updates meetings.
meeting:delete:meeting	Deletes meetings.
meeting:write:registrant	Adds meeting registrants.
cloud_recording:read:list_user_recordings	Lists user recordings.
cloud_recording:read:list_recording_files	Lists recording files.
cloud_recording:delete:meeting_recording	Deletes meeting recordings.
webinar:read:webinar	Reads webinar details.
webinar:read:list_webinars	Lists webinars.
webinar:read:list_past_participants	Lists past webinar participants.
webinar:read:list_registrants	Lists webinar registrants.
webinar:write:registrant	Adds webinar registrants.
user:read:user	Reads user information.
user:read:settings	Reads user settings.
user:read:list_collaboration_devices	Lists collaboration devices.
contact:read:list_contacts	Lists contacts.

Note

Additional scopes for whiteboard and marketplace features are also requested. The scopes listed above cover the primary meeting, recording, webinar, user, and contact actions.

Setting up the connector in Amazon Quick

Connect from the Available tab

If you want to use Default OAuth app authentication, you can connect directly from the **Available** tab without additional configuration.

1. In the Amazon Quick console, choose **Connectors**.
2. On the **Available** tab, find **Zoom** and choose **Connect**.
3. Complete the Zoom sign-in flow and grant the requested permissions.

To configure a connector with Custom OAuth app instead, use the **Create for your team** tab as described below.

Create from the Create for your team tab

After you complete any required Zoom configuration, create the connector in Amazon Quick.

1. In the Amazon Quick console, choose **Connectors**.
2. Choose the **Create for your team** tab.
3. Find and choose **Zoom**.

Note

If a Zoom connector already exists, a dialog appears with your existing connectors. To use an existing connector, choose it. To create a new one, choose **No, create new**.

4. Enter a **Name** for your connector. Optionally, choose **+ Add Description** to add a description.
5. For **Connection type**, choose **Public network**.
6. For **OAuth Configuration**, choose one of the following authentication methods and configure the required fields.

a. For **Default OAuth app**:

No additional credentials are needed. Choose **Next** to continue.

b. For **Custom OAuth app**, configure the following fields:

- **Base URL** (Optional) – The Zoom API base URL. Example: `https://api.zoom.us`
- **Client ID** – The client ID from your Zoom app.
- **Client secret** – The client secret from your Zoom app.
- **Token URL** – The token endpoint. Example: `https://zoom.us/oauth/token`
- **Authorization URL** – The authorization endpoint. Example: `https://zoom.us/oauth/authorize`
- **Redirect URL** – Pre-filled with the Amazon Quick callback URL.

7. Choose **Next**.

8. If you chose **Default OAuth app** or **Custom OAuth app**, a Zoom authorization window opens. Review the requested permissions and choose **Allow**.

9. On the **Review** page, review the available actions for the connector. Choose **Next**.

10. On the **Publish** page, choose who can access the connector. You can enable access for everyone in your organization or search for specific teams or groups.

11. Choose **Publish**.

Available actions

After you set up the connector, the following actions are available.

Zoom available actions

Category	Action	Description
Meetings	List Meetings	Lists meetings for a user.
Meetings	Get Meeting	Retrieves details for a specific meeting.
Meetings	Create Meeting	Creates a new meeting.
Meetings	Update Meeting	Updates an existing meeting.

Category	Action	Description
Meetings	Get Meeting Summary	Retrieves the AI-generated summary for a meeting.
Meetings	List Meeting Summary Templates	Lists available meeting summary templates .
Meetings	Get Past Meeting Participants	Retrieves participants from a past meeting.
Meetings	List Past Meeting Instances	Lists instances of a recurring meeting.
Meetings	Create Meeting Registrant	Adds a registrant to a meeting.
Recordings	List Recordings	Lists all recordings for a user.
Recordings	Get Meeting Recordings	Retrieves recordings for a specific meeting.
Recordings	List Archived Files	Lists archived recording files.
Webinars	List Webinars	Lists webinars for a user.
Webinars	Get Webinar	Retrieves details for a specific webinar.
Webinars	List Webinar Participants	Lists participants for a webinar.
Contacts	Search Company Contacts	Searches for contacts in the company directory.
Devices	List Devices	Lists Zoom devices.
Reports	Get Daily Usage Report	Retrieves a daily usage report.
Users	Get User	Retrieves information about the authenticated user.

Note

The actions that you can use depend on the permissions configured for your Zoom account and the data accessible to the authenticated user.

Managing and troubleshooting

To edit, share, or delete your connector, see [the section called “Managing existing integrations”](#).

Authentication issues

- **Sign-in fails (Default OAuth app or Custom OAuth app)** – Verify that your Zoom account is active and that you can sign in to zoom.us on the Zoom website directly. For Custom OAuth app, confirm that the redirect URI in your Zoom app matches the Amazon Quick callback URL.
- **Invalid client credentials (Custom OAuth app)** – Verify that the Client ID and Client secret match the values in your Zoom app. You can view your credentials from the app settings in the Zoom App Marketplace.
- **Insufficient permissions** – Verify that the scopes configured for your Zoom app include the permissions required for the actions that you want to use. See [the section called “Recommended scopes”](#).

ZoomInfo integration

With the ZoomInfo connector, you can look up company and contact data, and query firmographic and intent data directly in Amazon Quick through natural language.

Amazon Quick supports multiple authentication methods for ZoomInfo. Choose the method that best fits your organization's security requirements.

- **Default OAuth app** – Uses an AWS-managed OAuth application. No additional credentials are needed. Users authenticate directly with their ZoomInfo account.
- **Custom OAuth app** – Uses a customer-managed OAuth application. This option gives your organization full control over the OAuth configuration.
- **Service-to-Service OAuth** – Uses client credentials for server-to-server authentication without user interaction. Suitable for automated workflows.
- **API Key** – Uses an API key for authentication.

For more information about the authentication methods that Amazon Quick supports, see [the section called "Authentication methods"](#).

Before you begin

Make sure that you have the following before you set up the integration.

- An active ZoomInfo account with access to the data products you want to query.
- For **Custom OAuth app** or **Service-to-Service OAuth**: OAuth credentials from your ZoomInfo administrator. ZoomInfo uses Okta for OAuth.
- For **API Key**: A ZoomInfo-issued API key.
- For Amazon Quick subscription requirements, see [the section called "Set up integrations in the console"](#).

Configuring ZoomInfo

If you are using **Default OAuth app** authentication, skip this section and proceed to [the section called "Setting up the connector in Amazon Quick"](#).

For Custom OAuth app or Service-to-Service OAuth authentication, work with your ZoomInfo administrator to register an OAuth client in the ZoomInfo Okta tenant and add the Amazon Quick callback URL `https://{region}.quicksight.aws.amazon.com/sn/oauthcallback` as a redirect URI. Replace `{region}` with your AWS Region (for example, `us-east-1`). For step-by-step instructions, see [ZoomInfo MCP Overview](#) and [Developer Portal Guide](#) in the ZoomInfo documentation. Record the Client ID and Client Secret — you need them when you configure Amazon Quick. For API Key authentication, generate an API key from your ZoomInfo account.

Setting up the connector in Amazon Quick

Connect from the Available tab

If you want to use Default OAuth app authentication, you can connect directly from the **Available** tab without additional configuration.

1. In the Amazon Quick console, choose **Connectors**.
2. On the **Available** tab, find **ZoomInfo** and choose **Connect**.
3. Complete the ZoomInfo sign-in flow and grant the requested permissions.

To configure a connector with one of the other authentication methods, use the **Create for your team** tab as described below.

Create from the Create for your team tab

1. In the Amazon Quick console, choose **Connectors**.
2. Choose the **Create for your team** tab.
3. Find and choose **ZoomInfo**.
4. Enter a **Name** for your connector. Optionally, choose **+ Add Description** to add a description.
5. For **Connection type**, choose **Public network**.
6. For **OAuth Configuration**, choose one of the following authentication methods and configure the required fields.

a. For **Default OAuth app**:

No additional credentials are needed. Choose **Next** to continue.

b. For **Custom OAuth app**, configure the following fields:

- **Client ID** – The client ID from your ZoomInfo OAuth client.
- **Public OAuth client** (Optional) – Select this option if your ZoomInfo OAuth client is configured as a public client (no client secret).
- **Client secret** – The client secret from your ZoomInfo OAuth client.
- **Token URL** – The token endpoint. Default: `https://okta-login.zoominfo.com/oauth2/default/v1/token`
- **Authorization URL** – The authorization endpoint. Default: `https://okta-login.zoominfo.com/oauth2/default/v1/authorize`
- **Redirect URL** – Pre-filled with the Amazon Quick callback URL.

c. For **Service-to-Service OAuth**, configure the following fields:

- **Client ID** – The client ID from your ZoomInfo OAuth client.
- **Client secret** – The client secret from your ZoomInfo OAuth client.
- **Token URL** – The token endpoint. Default: `https://okta-login.zoominfo.com/oauth2/default/v1/token`

d. For **API Key**, configure the following fields:

- **API Key** – The ZoomInfo API key.

- **Email** (Optional) – The email address associated with the API key.
7. Choose **Next**.
 8. If you chose **Default OAuth app** or **Custom OAuth app**, a ZoomInfo authorization window opens. Review the requested permissions and choose **Allow**.
 9. On the **Review** page, review the available actions for the connector. Choose **Next**.
 10. On the **Publish** page, choose who can access the connector. You can enable access for everyone in your organization or search for specific teams or groups.
 11. Choose **Publish**.

Available actions

After you set up the connector, the actions exposed by ZoomInfo are available. To see the current set of actions for your connector, go to the connector's **Available actions** view in the Amazon Quick console.

Managing and troubleshooting

To edit, share, or delete your connector, see [the section called “Managing existing integrations”](#).

Authentication issues

- **Sign-in fails (Default OAuth app or Custom OAuth app)** – Verify that your ZoomInfo account is active. For Custom OAuth app, confirm that the redirect URI in your ZoomInfo OAuth client matches the Amazon Quick callback URL.
- **Invalid client credentials (Custom OAuth app or Service-to-Service OAuth)** – Verify that the Client ID and Client secret match the values in your ZoomInfo OAuth client.
- **API Key authentication fails** – Verify that the API key has not been revoked.

Support escalation path

If you cannot resolve an issue for an integration using the troubleshooting steps, follow this escalation path:

1. Gather diagnostic information:
 - Integration configuration details

- Error messages and timestamps
 - Steps to reproduce the issue
 - Environment details if applicable (Online vs. Server, version)
2. Check Amazon Quick service health and known issues in the console.
 3. Contact AWS Support through the Amazon Quick console or AWS Support Center.
 4. For complex authentication or permission issues, be prepared to involve your administrator.

Bring Your Own Amazon Q Business Index (BYOI)

Amazon Quick enables you to use your existing Amazon Q Business indexes as data sources. You can leverage your enterprise data without having to recreate indexes. This capability, known as Bring Your Own Index (BYOI), lets you connect your Amazon Q Business indexes to Amazon Quick and use them alongside other data sources for comprehensive analytics and intelligent responses.

BYOI supports two implementation methods:

IDC Implementation

Uses IAM Identity Center for authentication. Requires both Amazon Q Business and Amazon Quick to authenticate through IAM Identity Center in the same AWS account and region.

Non-IDC Implementation

Supports multiple authentication methods including native identities, AWS Managed Microsoft AD, and IAM federation. All Amazon Quick users automatically receive access to connected Amazon Q Business indexes.

Topics

- [Overview of Amazon Q Business indexes in Amazon Quick](#)
- [Prerequisites](#)
- [Supported authentication methods](#)
- [Setting up permissions](#)
- [Creating knowledge bases from Amazon Q Business indexes](#)
- [Sharing Amazon Q Business index knowledge bases](#)
- [Using Amazon Q Business index knowledge bases](#)

- [Limitations](#)
- [Billing](#)
- [Feature comparison](#)
- [Troubleshooting](#)
- [Security best practices](#)
- [User types and capabilities](#)
- [Common use cases](#)
- [Interacting with your Amazon Q Business indexes](#)

Overview of Amazon Q Business indexes in Amazon Quick

Amazon Q Business indexes contain indexed enterprise data that you can now use directly in Amazon Quick. When you connect an Amazon Q Business index to Amazon Quick, it becomes available as a knowledge base. You can use it in spaces, agents, and automations, just like any other knowledge base in Amazon Quick.

Key benefits of using Amazon Q Business indexes in Amazon Quick include:

Leverage existing data investments

Use your existing Amazon Q Business indexes and indexed data directly in Amazon Quick. You don't need to recreate indexes or upload the same data again.

Unified analytics experience

Query across multiple data sources including Amazon Q Business indexes, uploaded documents, and structured data. Use Amazon Quick's unified analytical environment for all your data.

Consistent security and permissions

Amazon Q Business index knowledge bases in Amazon Quick maintain the same security and access controls as in Amazon Q Business. Users only see content they have permission to access.

Enhanced collaboration

Share Amazon Q Business index knowledge bases with other users in Amazon Quick, add them to spaces, and use them with agents and automations to create comprehensive business solutions.

Prerequisites

Before you can use Amazon Q Business indexes in Amazon Quick, ensure that you meet the following prerequisites:

Common Prerequisites

- You have an existing Amazon Q Business index with indexed data.
- Both the Amazon Q Business index and the Amazon Quick instance are in the same AWS account and region.
- You have administrator permissions in Amazon Quick.

IDC Implementation Prerequisites

- AWS Identity Center is enabled and configured.
- Both Amazon Q Business and Amazon Quick authenticate through IAM Identity Center.
- The IAM Identity Center region and Amazon Q Business index region are the same.
- You have access to both the Amazon Q Business index and AWS Identity Center administration.

Supported authentication methods

The authentication methods supported depend on your implementation type:

IDC Implementation

- Amazon Quick: AWS Identity Center authentication only
- Amazon Q Business: AWS_IAM_IDC

Non-IDC Implementation

- Amazon Quick:
 - Native identities (username/password)
 - AWS Managed Microsoft AD
 - IAM federation
- Amazon Q Business: AWS_QUICKSIGHT_IDP

Setting up permissions

To use Amazon Q Business indexes in Amazon Quick, you need to set up the appropriate permissions based on your implementation method:

Initial Setup

1. Sign in to the Amazon Quick console as an administrator.
2. Navigate to the **Admin** section.
3. Select **AWS Resources**.
4. Choose **Amazon Q Business** from the list of available data sources.
5. Choose **Select Applications**.

Application Setup

You can either connect to an existing Amazon Q Business application or create a new one:

1. Choose one of the following options:
 - **Connect to existing Amazon Q Business application** - Select an existing application from your account.
 - **Create new Amazon Q Business application** - Create a new application. The new application will use the same authentication used by your Amazon Quick instance setup.
2. For new applications, the system automatically configures authentication based on your Amazon Quick instance setup.
3. Wait for application creation to complete.
4. You will be redirected to the Amazon Q Business application to configure indexes and data sources.

Access Management by Implementation

IDC Implementation

- Access is managed through AWS Identity Center.
- Access to the Amazon Q Business application is managed through the Amazon Q Business console.

Non-IDC Implementation

- All Amazon Quick users automatically gain access to connected Amazon Q Business indexes.
- No additional access management required in Amazon Q Business.

Once permissions are set up, you can use your Amazon Q Business index as a knowledge base in Amazon Quick, and Admin users can create knowledge bases from Amazon Q Business indexes.

Creating knowledge bases from Amazon Q Business indexes

After setting up permissions, Admin or Admin Pro users can create knowledge bases from Amazon Q Business indexes:

1. Sign in to the Amazon Quick console as an Admin or Admin pro user.
2. Navigate to **Knowledge bases**.
3. Choose **Create knowledge base**.
4. Select **Amazon Q Business** as the data source.
5. Choose the Amazon Q Business index you want to use from the list of available indexes.
6. Provide a name and description for the knowledge base.
7. Choose **Create** to create the knowledge base.

Note

If multiple Admin users create knowledge bases from the same Amazon Q Business index, these knowledge bases will functionally be identical.

Sharing Amazon Q Business index knowledge bases

Users who have access to an Amazon Q Business index in Amazon Q Business will automatically have the corresponding knowledge base in Amazon Quick shared with them when it is created. The knowledge base can be shared with other users after creation, but only the user who created the knowledge base can share it with other users.

In an IDC implementation, any changes to the permissions in Amazon Q Business application after the creation of knowledge base will not take effect in Amazon Quick.

If you are an admin and want to share the knowledge base with other users in Amazon Quick:

1. Navigate to **Knowledge bases**.
2. Select the Amazon Q Business index knowledge base you want to share.
3. Choose **Share**.
4. Select the users or groups you want to share the knowledge base with.
5. Choose **Share** to confirm.

Using Amazon Q Business index knowledge bases

Once created, Amazon Q Business index knowledge bases can be used in Amazon Quick just like any other knowledge base:

Using in spaces

Admins can add Amazon Q Business index knowledge bases to spaces:

1. Navigate to the space where you want to add the knowledge base.
2. Choose **Add resources**.
3. Select **Knowledge bases**.
4. Choose the Amazon Q Business index knowledge base from the list.
5. Choose **Add** to confirm.

Using in output cards

Admins can use Amazon Q Business index knowledge bases via Spaces in output cards.

Using with agents

Admins can add Amazon Q Business index knowledge bases to custom agents:

1. Navigate to **Agents**.
2. Select an existing agent or create a new one.
3. In the agent configuration, choose **Add knowledge bases**.
4. Select the Amazon Q Business index knowledge base from the list.
5. Choose **Add** to confirm.

Using with automations

Admins can add Amazon Q Business index knowledge bases to automations:

1. Navigate to **Automations**.
2. Select an existing automation or create a new one.
3. In the automation configuration, add a step that uses knowledge bases.
4. Select the Amazon Q Business index knowledge base from the list.
5. Configure the step and save the automation.

Querying knowledge bases

Readers can query Amazon Q Business index knowledge bases through the Amazon Quick web application. However, a user will only be able to get a response from the Amazon Q Business index if they also have access to the Amazon Q Business application. To query the knowledge base:

1. Navigate to the Amazon Quick web application.
2. Select a space that contains the Amazon Q Business index knowledge base or use the default agent.
3. Enter your query in the chat interface.
4. View the response, which includes citations and clickable links to the source documents from the Amazon Q Business index.

Limitations

Note

In an IDC implementation, when a Amazon Q Business knowledge base is first created in Amazon Quick, access to the knowledge base is automatically granted to users with access to the selected Amazon Q Business index. For additional users to have access to a knowledge base, the admin must configure user access in both the Amazon Q Business console and Amazon Quick knowledge base permissions pages.

When using Amazon Q Business indexes in Amazon Quick, be aware of the following limitations:

General Limitations

- Amazon Q Business index knowledge bases cannot be changed like other knowledge bases in Amazon Quick.
- Amazon Q Business index knowledge bases only support the docs types supported by Amazon Q Business.
- QApps, Actions, and Amazon Q Business chat guardrails are not included in the BYOI capability.
- Amazon Q Business indexes must be in the same AWS account and region as Amazon Quick.

IDC Implementation Limitations

- Both Amazon Quick and Amazon Q Business must use the same instance of IAM Identity Center.

Index Quotas

- You can connect up to two Amazon Q Business indexes per Region to Amazon Quick in the current release.
- This quota cannot be increased.
- Once indexes are selected and saved in a Amazon Quick instance, they cannot be directly unselected.

Billing

When using Amazon Q Business indexes in Amazon Quick, billing works as follows:

IDC Implementation Billing

- You are billed for index and subscription costs in Amazon Q Business (recommended at least 1 user).
- You are billed for subscription costs in Amazon Quick based on your Amazon Quick user subscriptions.
- There are no additional charges for connecting Amazon Q Business indexes to Amazon Quick.

Non-IDC Implementation Billing

- No explicit Amazon Q Business subscription charges apply in this implementation model.
- You are billed for subscription costs in Amazon Quick based on your Amazon Quick user subscriptions.

For more information about Amazon Quick pricing, see [Amazon Quick pricing](#).

Feature comparison

The following table compares key features between IDC and Amazon Q Business implementations:

Feature	IDC Implementation	Non-IDC Implementation
User management	AWS Identity Center	Amazon Quick
Amazon Quick authentication methods	Identity Center Only	Native identities (username /password), AWS Managed Microsoft AD, IAM federation
Amazon Q Business authentication methods	AWS_IAM_IDC	AWS_QUICKSIGHT_IDP
Share permissions	Amazon Q Business Console and Knowledge base permissions page	Amazon Quick Knowledge Base Permission page (Automatic)
Index compatibility	All indexes	All indexes

Troubleshooting

Amazon Q Business not seen in Integrations page

Symptoms

- Amazon Q Business option missing from Integrations page
- Cannot create new Amazon Q Business integration

Resolution

- Only Admin users have access to create a Amazon Q Business/BYOI knowledge base
- Verify user has Admin persona permissions

Failed to fetch Amazon Q Business applications

Resolution

- Confirm Amazon Q Business is enabled in the Admin console
- Try logging out and back in to refresh the session, then retry the operation

Amazon Q Business application not seen in the list of applications displayed during knowledge base creation

Symptoms

- Amazon Q Business applications list is empty on Create Knowledge Base page
- Amazon Q Business applications list is populated but missing expected applications

Resolution

- Check if missing Amazon Q Business applications were granted permissions in the AWS resources page of the Admin console

Failed to create dataset. Chat instance is not ready. Please try again later

Symptoms

- Knowledge base creation fails with error "Chat instance is not ready. Please try again later"
- Unable to complete knowledge base creation process

Resolution

- If this is the first time creating a knowledge base in Amazon Quick, wait 5 minutes and retry the operation

Security best practices

- Regularly review access permissions
- Monitor user activities
- Implement least-privilege access
- Maintain authentication method security

User types and capabilities

Different user types have different capabilities when working with Amazon Q Business indexes in Amazon Quick:

Amazon Q Business Administrator

Manages Amazon Q Business indexes in the Amazon Q Business console.

Amazon Quick Administrator

Enables permissions for Amazon Quick to access Amazon Q Business indexes in the same AWS account and region.

Amazon Quick Admin

Creates knowledge bases from Amazon Q Business indexes, shares them with other users, and adds them to spaces.

Amazon Quick Admin

Creates agents and automations that leverage Amazon Q Business index knowledge bases.

Amazon Quick Reader

Uses agents and automations that leverage Amazon Q Business index knowledge bases, and queries these knowledge bases through the Amazon Quick web application.

Common use cases

Here are some common use cases for using Amazon Q Business indexes in Amazon Quick:

Finding relevant information

Query about specific topics across your enterprise data, such as HR policies, support documentation, or compliance guidelines.

Example: "What is the maternity leave policy for my company in Washington state?" where national guidelines are in an Amazon Q Business index and state-specific information is in Amazon Quick.

Business analytics and insights

Extract key findings from status updates across multiple groups or analyze the health of your product portfolio.

Example: "Analyze the health of my product portfolio" where product sales data is in an Amazon Q Business index, and product engineering and support data is in Amazon Quick.

Content creation and summarization

Create summaries or reports based on documents from multiple sources.

Example: "Summarize last week's Service WBR and Sales update into a one-pager reflecting key initiatives" where the Service WBR is indexed in an Amazon Q Business index and the latest sales update is uploaded to Amazon Quick.

Interacting with your Amazon Q Business indexes

Once you've connected your Amazon Q Business index to Amazon Quick, you can interact with it through chat using agents.

1. From the left navigation menu in the console, select **Agents**.
2. From the **Actions** column for the agent you want to share, select the menu icon, and then select **Chat**.
3. You can then narrow down the resource to just the Knowledge Base created from the Amazon Q Business index by selecting the All Resources option and clicking on your Amazon Q Business index from the Knowledge Bases sub-tab.

When you ask questions, the agent analyzes data from your Amazon Q Business index alongside other data sources to provide comprehensive responses.

You can ask complex questions in plain language about your enterprise data and get detailed responses that draw from your Amazon Q Business index's indexed content. The chat interface allows you to:

- Query your enterprise data using natural language
- Get contextual responses that combine data from your Amazon Q Business index with other Amazon Quick data sources
- View source citations with clickable references to see where information originated
- Continue conversations and build on previous queries

For detailed information about using the chat interface and working with agents, see [the section called "Use a chat agent"](#).

Bring your own Amazon Bedrock managed knowledge base

Amazon Quick can connect Amazon Bedrock managed knowledge bases as knowledge sources. After connection, a managed knowledge base is treated like any other knowledge base in Amazon Quick. You can add it to spaces and it is queried automatically during chat. You can connect up to 2 managed knowledge bases per Amazon Quick instance. The managed knowledge base can live in the same AWS account as Amazon Quick or in a different account.

Topics

- [Prerequisites](#)
- [Setting up permissions](#)
- [Creating knowledge bases from Amazon Bedrock managed knowledge bases](#)
- [Access control list \(ACL\) support](#)
- [Billing](#)
- [Limitations](#)
- [Security best practices](#)
- [Monitoring and observability](#)
- [Troubleshooting](#)

Prerequisites

- A Amazon Bedrock managed knowledge base that has been created and configured. For more information, see [Build a managed knowledge base](#) in the *Amazon Bedrock User Guide*.
- The managed knowledge base and the Amazon Quick instance are in the same AWS Region.
- Administrator-level permissions in Amazon Quick to manage AWS resource integrations.
- For cross-account setups: the ability to attach a resource policy to the knowledge base, or coordination with the knowledge base owner to attach the policy.

Setting up permissions

Amazon Quick authenticates to Amazon Bedrock using IAM SigV4 credentials via the Amazon Quick service role. The service role requires the following actions on the knowledge base:

- `bedrock:Retrieve`
- `bedrock:GetDocumentContent`

The following example shows the ARN format for Amazon Bedrock managed knowledge bases:

```
arn:aws:bedrock:REGION:ACCOUNT_ID:knowledge-base/KB_ID
```

Cross-account setup

Cross-account access is required when the managed knowledge base is owned by a different AWS account than the one running your Amazon Quick instance. For example, a central data team might own the knowledge base in one account while your Amazon Quick instance runs in a separate account.

In this case, you must attach a resource policy to the knowledge base in the owner's account before granting the service role access. The resource policy authorizes the Amazon Quick service role to call the knowledge base across account boundaries.

The following resource policy example grants the Amazon Quick service role permission to access your managed knowledge base. Replace the placeholder values with your specific account IDs, Region, and knowledge base ID:

```
{
```



```
"Version": "2012-10-17",
"Statement": [
  {
    "Effect": "Allow",
    "Principal": {
      "AWS": "arn:aws:iam::QUICK_ACCOUNT_ID:role/service-role/aws-quicksight-
service-role-v0"
    },
    "Action": [
      "bedrock:Retrieve",
      "bedrock:GetDocumentContent"
    ],
    "Resource": "arn:aws:bedrock:REGION:KB_OWNER_ACCOUNT_ID:knowledge-
base/KB_ID"
  }
]
```

See [Share a managed knowledge base across accounts](#) in the *Amazon Bedrock User Guide*.

Granting the service role access

For both same-account and cross-account setups, grant the Amazon Quick service role access to the managed knowledge base through the admin console:

1. Sign in to the Amazon Quick console as an administrator.
2. Navigate to **Manage account > AWS Resources > Bedrock**.
3. Add your managed knowledge base ARN to the Amazon Quick service role using the provided form. Amazon Quick attaches the necessary IAM policy to the service role, scoped to the entered ARN.

Creating knowledge bases from Amazon Bedrock managed knowledge bases

This topic shows you how to create a knowledge base in Amazon Quick from an existing Amazon Bedrock managed knowledge base. After completing the permission setup, administrators can use this procedure to integrate the managed knowledge base into Amazon Quick.

To create a knowledge base from an Amazon Bedrock managed knowledge base

1. Sign in to the Amazon Quick console as an administrator.
2. Navigate to **Knowledge bases > Set up new knowledge base**.
3. Select **Amazon Bedrock managed knowledge base** as the source type.
4. Complete the form and choose **Create**.

After creation, the managed knowledge base is available in Amazon Quick and can be added to spaces, agents, and automations.

Access control list (ACL) support

Amazon Bedrock managed knowledge bases with ACL-enabled data source connectors support document-level access control. When a user queries the knowledge base, Amazon Quick automatically passes the user's identity to Amazon Bedrock. Amazon Bedrock then filters retrieval results so that users only see documents they are authorized to access in the original data source.

No additional configuration is required in Amazon Quick to enable ACL support. The access control rules are configured on the Amazon Bedrock managed knowledge base data source connectors. See [Connect a data source](#) in the *Amazon Bedrock User Guide*.

Billing

The knowledge base owner is charged for Amazon Bedrock usage costs, including retrieval requests and data ingestion. Amazon Quick subscription costs are billed separately based on your Amazon Quick user subscriptions. There are no additional charges from Amazon Quick for connecting an Amazon Bedrock managed knowledge base.

Limitations

- This feature is available in the following AWS Regions: US East (N. Virginia), US West (Oregon), Europe (Ireland), and Asia Pacific (Sydney).
- You can connect up to 2 managed knowledge bases per Amazon Quick instance.
- The managed knowledge base and Amazon Quick instance must be in the same AWS Region.
- Only Amazon Bedrock managed knowledge bases are supported. Custom knowledge bases with self-managed vector stores are not supported.
- Cross-account access requires a resource policy on the knowledge base.

Security best practices

- Use cross-account resource policies with specific role ARNs rather than granting access to entire accounts.
- Regularly review the resource policies attached to your managed knowledge bases.
- Monitor access patterns and review audit logs regularly.

Monitoring and observability

Amazon Bedrock emits AWS CloudTrail data events for Retrieve and GetDocumentContent API calls on your managed knowledge base. You can use these events to monitor successful usage and detect access failures. The following walkthrough sets up a CloudTrail trail that delivers data events to a Amazon CloudWatch Logs log group, then creates metric filters and alarms to alert on errors.

Step 1: Enable CloudTrail to Amazon CloudWatch for Amazon Bedrock knowledge base events

Set up a CloudTrail trail that captures Amazon Bedrock knowledge base data events and delivers them to a Amazon CloudWatch Logs log group. Create the trail and log group in the same Region as your managed knowledge base. You can also do this through the CloudTrail console, which creates the log group and IAM role automatically.

Create the S3 bucket, apply a bucket policy that allows CloudTrail to write, and create the Amazon CloudWatch Logs log group:

```
aws s3 mb s3://YOUR_TRAIL_BUCKET

aws s3api put-bucket-policy --bucket YOUR_TRAIL_BUCKET --policy '{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {"Service": "cloudtrail.amazonaws.com"},
      "Action": "s3:GetBucketAcl",
      "Resource": "arn:aws:s3:::YOUR_TRAIL_BUCKET"
    },
    {
      "Effect": "Allow",
      "Principal": {"Service": "cloudtrail.amazonaws.com"},
```

```

    "Action": "s3:PutObject",
    "Resource": "arn:aws:s3:::YOUR_TRAIL_BUCKET/AWSLogs/ACCOUNT_ID/*",
    "Condition": {"StringEquals": {"s3:x-amz-acl": "bucket-owner-full-control"}}
  }
]
}'

```

```
aws logs create-log-group --log-group-name bedrock-kb-cloudtrail
```

Create an IAM role that allows CloudTrail to deliver logs to the log group:

```

aws iam create-role \
  --role-name CloudTrail-CWLogs-Role \
  --assume-role-policy-document '{
    "Version": "2012-10-17",
    "Statement": [{
      "Effect": "Allow",
      "Principal": {"Service": "cloudtrail.amazonaws.com"},
      "Action": "sts:AssumeRole"
    }]
  }'

aws iam put-role-policy \
  --role-name CloudTrail-CWLogs-Role \
  --policy-name CWLogsPolicy \
  --policy-document '{
    "Version": "2012-10-17",
    "Statement": [{
      "Effect": "Allow",
      "Action": ["logs:CreateLogStream", "logs:PutLogEvents"],
      "Resource": "arn:aws:logs:REGION:ACCOUNT_ID:log-group:bedrock-kb-cloudtrail:*"
    }]
  }'

```

Create the trail, configure it to capture Amazon Bedrock knowledge base data events, and start logging:

```

aws cloudtrail create-trail \
  --name bedrock-kb-data-events \
  --s3-bucket-name YOUR_TRAIL_BUCKET \
  --cloud-watch-logs-log-group-arn "arn:aws:logs:REGION:ACCOUNT_ID:log-group:bedrock-kb-cloudtrail:*" \
  --cloud-watch-logs-role-arn "arn:aws:iam::ACCOUNT_ID:role/CloudTrail-CWLogs-Role"

```

```
aws cloudtrail put-event-selectors \
  --trail-name bedrock-kb-data-events \
  --advanced-event-selectors '[{
    "Name": "BedrockKBDataEvents",
    "FieldSelectors": [
      {"Field": "eventCategory", "Equals": ["Data"]},
      {"Field": "resources.type", "Equals": ["AWS::Bedrock::KnowledgeBase"]}
    ]
  }]'

aws cloudtrail start-logging --name bedrock-kb-data-events
```

Step 2: Create metric filters for Amazon Quick retrieval calls

Create metric filters to track retrieval calls and errors when Amazon Quick invokes your managed knowledge base.

Call volume metric — tracks retrieval call volume to your managed knowledge base.

```
aws logs put-metric-filter \
  --log-group-name bedrock-kb-cloudtrail \
  --filter-name QuickSight-Retrieve-CallCount-KB_ID \
  --filter-pattern '{ $.eventSource = "bedrock.amazonaws.com" && $.eventName
= "Retrieve" && $.userIdentity.invokedBy = "quicksight.amazonaws.com" &&
$.resources[0].ARN = "*KB_ID*" }' \
  --metric-transformations '[{
    "metricName": "RetrieveCallCount",
    "metricNamespace": "QuickSight/BedrockKB",
    "metricValue": "1",
    "dimensions": {"KnowledgeBaseId": "$.resources[0].ARN"}
  }]'
```

Error metric — the following example counts Retrieve calls from Amazon Quick that fail with `AccessDenied`. This alarm fires when the Amazon Quick service role is missing permissions or when a cross-account resource policy is misconfigured. You can create additional metric filters for other error codes such as `ThrottlingException` depending on your monitoring needs.

```
aws logs put-metric-filter \
  --log-group-name bedrock-kb-cloudtrail \
  --filter-name QuickSight-Retrieve-AccessDenied-KB_ID \
```

```
--filter-pattern '{ $.eventSource = "bedrock.amazonaws.com" && $.eventName
= "Retrieve" && $.userIdentity.invokedBy = "quicksight.amazonaws.com" &&
$.resources[0].ARN = "*KB_ID*" && $.errorCode = "AccessDenied" }' \
--metric-transformations '[{
  "metricName": "RetrieveAccessDeniedCount",
  "metricNamespace": "QuickSight/BedrockKB",
  "metricValue": "1",
  "dimensions": {"KnowledgeBaseId": "$.resources[0].ARN"}
}]'
```

Step 3: Create a CloudWatch alarm

Create an alarm that triggers when access denied errors occur. Optionally add an SNS topic ARN to `--alarm-actions` for notifications.

```
aws cloudwatch put-metric-alarm \
  --alarm-name QuickSight-Retrieve-AccessDenied-KB_ID \
  --metric-name RetrieveAccessDeniedCount \
  --namespace "QuickSight/BedrockKB" \
  --dimensions Name=KnowledgeBaseId,Value="arn:aws:bedrock:REGION:ACCOUNT_ID:knowledge-
base/KB_ID" \
  --statistic Sum \
  --period 300 \
  --threshold 1 \
  --comparison-operator GreaterThanOrEqualToThreshold \
  --evaluation-periods 1 \
  --treat-missing-data notBreaching
```

Troubleshooting

Troubleshoot access denied errors

When you query an Amazon Bedrock managed knowledge base in Amazon Quick, retrieval is authorized only when every permission layer allows it. If you receive the error *Insufficient permissions to access this knowledge base*, work through the following layers in order and stop at the first layer that is not satisfied.

Tip

First, check AWS CloudTrail for the AccessDenied event to identify which principal and action were rejected. Then go directly to the matching layer in the following list.

1. **Service role permissions.** Confirm that the Amazon Quick service role grants `bedrock:Retrieve` and `bedrock:GetDocumentContent` on the knowledge base.
2. **Cross-account resource policy.** For cross-account setups, confirm that the knowledge base resource policy grants those same actions to the Amazon Quick service role.
3. **Customer managed key.** If the knowledge base is encrypted with a customer managed key, confirm that the service role has `kms:Decrypt` and `kms:DescribeKey` on that key.
4. **Knowledge base ARN.** Confirm that the knowledge base ARN is correct in the Amazon Quick admin configuration. Amazon Quick scopes the service role policy to the ARN you enter, so an incorrect ARN denies retrieval.
5. **IAM policy assignments.** If IAM policy assignments are enabled in the account, Amazon Quick applies a per-user session policy at retrieval time and intersects it with the service role permissions. An action succeeds only if both the service role and the user's assigned IAM policy allow it. Check this layer when the knowledge base works for an administrator but fails for a specific user or group.

To grant retrieval permissions to affected users

1. Open the [Quick console](#).
2. Choose **Manage account**, and then choose **IAM policy assignments**.
3. Create a new assignment, or edit the IAM policy behind an existing assignment, to allow the knowledge base retrieval actions. For example:

```
{
  "Sid": "AllowQuickFMKBRetrieve",
  "Effect": "Allow",
  "Action": ["bedrock:Retrieve", "bedrock:GetDocumentContent"],
  "Resource": "arn:aws:bedrock:REGION:ACCOUNT_ID:knowledge-base/KB_ID"
}
```

4. Add the affected users or groups to the assignment, and then save and enable it.

For more information about IAM policy assignments, see [the section called “IAM policy assignments”](#).

Connect Amazon Quick to your existing tools with extensions

Note

The Amazon Quick browser extension, Microsoft Outlook extension, Microsoft Word extension, Microsoft Excel extension, and Microsoft PowerPoint extension don't require admin setup to install and use. For the Slack and Microsoft Teams extensions, your Quick admin must [configure Amazon Quick access](#) before authors can deploy them.

Amazon Quick extensions enable you to optimize your daily tasks with assistance from an AI tool added to the tools you already use. Whether you're browsing the web, collaborating in team channels, managing emails, or creating documents, Quick extensions bring AI-powered assistance directly into your existing workflows without requiring you to switch between applications.

The following extension types are available:

- **Browser Extension** - Enhance your web browsing experience with AI-powered page summarization, content analysis, and knowledge access across Google Chrome, Firefox, and Microsoft Edge browsers. No admin setup required.
- **Microsoft Outlook extension** - Manage your entire inbox using natural language. Summarize unread messages, organize your inbox, schedule meetings, and draft contextual responses using your Quick enterprise data. No admin setup required.
- **Microsoft Excel extension** - Perform in-depth analysis of your workbooks, import and clean up data, generate tables and charts, and automate spreadsheet operations using natural language. No admin setup required.
- **Microsoft Word extension** - Generate formatted documents, make edits while maintaining writing style and structure, and redline documents with track changes and comments. No admin setup required.
- **Microsoft PowerPoint extension** - Create visually appealing presentations, edit slides using natural language, and streamline operations such as restructuring presentations and applying transitions. No admin setup required.

- **Slack extension** - Access Quick capabilities directly within Slack channels and direct messages, enabling seamless collaboration and knowledge sharing without leaving your team communication platform. Requires admin to configure extension access.
- **Microsoft Teams extension** - Integrate Quick into your Teams workspace for real-time collaboration, thread analysis, and knowledge retrieval during meetings and conversations. Requires admin to configure extension access.

Each extension maintains enterprise-grade security and privacy standards while providing contextual AI assistance tailored to the specific application environment. All extensions integrate with your organization's Quick knowledge bases, custom agents, and configured data sources to provide relevant, accurate responses based on your permissions and access levels.

Topics

- [Amazon Quick user interaction with extensions](#)
- [Amazon Quick browser extension](#)
- [Amazon Quick Microsoft Outlook extension](#)
- [Amazon Quick Microsoft Excel extension](#)
- [Amazon Quick Microsoft Word extension](#)
- [Amazon Quick Microsoft PowerPoint extension](#)
- [Amazon Quick Slack extension](#)
- [Amazon Quick Microsoft Teams extension](#)

Amazon Quick user interaction with extensions

Amazon Quick extensions enable different user roles to collaborate effectively in creating, managing, and utilizing AI-powered analytics tools. Each role—administrators, authors, and end users—has distinct responsibilities and capabilities when working with extensions. Understanding these role-based interactions ensures proper governance, security, and optimal user experience across your organization's Amazon Quick deployment.

Topics

- [User capabilities and permissions](#)
- [Amazon Quick administrator responsibilities](#)
- [Amazon Quick author capabilities](#)

- [Amazon Quick end user access](#)
- [Extension access troubleshooting](#)

User capabilities and permissions

Amazon Quick implements a hierarchical permission model for extensions that balances organizational control with user productivity. Administrators establish the foundational policies and infrastructure, authors manage the deployment and configuration of specific extension instances, and readers benefit from the AI-powered assistance within their authorized scope. The following table shows how each user type in Amazon Quick interacts with extensions.

User Type	Capabilities
Administrators	• Enable/remove extension access to their organization's workspaces and organizations • Create extensions for Slack, Teams, Word, and Outlook • Assign owners for extensions created • Control whether extensions are available for users in Amazon Quick • Use all enabled extensions
Authors	• Create and install extensions using administrator enabled extension access • Edit/delete extensions for which they are owners • Use all enabled extensions
Readers (Reader Pro only)	• Use all enabled extensions

Beyond the broad user type categories, Amazon Quick implements granular permission controls that determine specific actions users can perform with individual extensions. These permissions operate independently of user types, allowing administrators to fine-tune access based on organizational needs and security requirements. The following table outlines how user permissions determine what you can do with a extension:

Permissions Type	Permissions
Owners	• Edit extension name and description • Install the extension in a workspace or organization • Grant ownership and editing permissions to specific users and groups
Viewers	• View extension name and description • Access links to open in Slack, M365, and browser • Use the extensions

Important

The Amazon Quick browser extension extension doesn't need any admin setup to install and use.

These permission structures ensure that extension access aligns with your organization's security requirements while enabling appropriate levels of functionality for each user role. Administrators maintain control over extension availability and initial setup, authors handle deployment and configuration, and all users can benefit from the AI-powered assistance once extensions are properly configured.

Amazon Quick administrator responsibilities

Amazon Quick administrators, or users granted admin privileges, establish the foundational infrastructure that enables extension functionality across the organization. They handle the critical setup and security configurations that allow Amazon Quick authors to deploy extensions and end users to access them safely.

Note

Creating and adjusting extension access requires IAM administrator privileges. All administrators will see the extension links, but if not [signed in with IAM credentials](#), you will need to sign in with appropriate IAM permissions to manage extension access. This must be an IAM user who is an Amazon Quick administrator and not non-IAM admin users.

Responsibility area	Administrator capabilities
Infrastructure setup	• Connect extension applications to organizational systems (Slack workspaces, Microsoft 365 tenants) • Handle OAuth authorization and tenant-level permissions • Manage platform-specific installation methods: manifest-only for Word/Outlook, OAuth-only for Slack/Teams
Security and compliance	• Manage enterprise security and compliance settings • Control which extension types are available to the organization

Amazon Quick author capabilities

Authors bridge the gap between administrative infrastructure and end-user access by creating, configuring, and deploying specific extension instances. They manage the day-to-day operational aspects of extension deployment and maintenance.

Author capabilities depend on the permission level granted by administrators. There are two permission types that determine what authors can do:

- **Viewers:** Can view extension name and description, access links to open extensions, and use the extensions.
- **Owners:** Can edit extension name and description, install extensions in workspaces or organizations, and share ownership permissions with users and groups.

This translates to two main setup flows:

- **Limited permissions (view, share, delete only):** Authors can only use the basic landing page functionality to view, share, and delete extensions after admin completes all setup.
- **Full permissions (deploy, view, share, delete, edit):** Authors can download manifests or use OAuth deployments to complete app setup, rename extensions, and access editing features.

Capability area	Author functions
Configuration and deployment	• Download and manage installation manifests for deployment (with full permissions) • Deploy extension instances using OAuth or manifest-based methods (with full permissions) • Complete app setup and installation (with full permissions)
Management and sharing	• Manage extension sharing and access permissions • View, share, and delete deployed extensions • Edit extensions and access creation features (with full permissions)

Amazon Quick end user access

End users represent the primary beneficiaries of extension functionality, accessing AI-powered assistance directly within their existing workflow applications. Their interaction with extensions focuses on daily productivity and seamless integration with organizational knowledge.

Access level	User capabilities
Daily usage and productivity	• Access extensions directly within existing workflow applications • Use extensions for knowledge access, document analysis, and action taking • Benefit from seamless integration without context switching between applications
Knowledge and actions	• Access organizational knowledge bases and custom agents through extension interfaces • Perform external actions in third-party applications using configured action connectors • Receive permissions-aware responses based on individual access levels

The following table shows how each user type in Amazon Quick interacts with extensions.

User Type	Capabilities
Administrators	• Enable/remove extension access to their organization's workspaces and organizations • Create extensions for Slack, Teams, Word, and Outlook • Assign owners for extensions created • Control whether extensions are available for users in Amazon Quick • Use all enabled extensions
Authors	• Create and install extensions using administrator enabled extension access • Edit/delete extensions for which they are owners • Use all enabled extensions
Readers (Reader Pro only)	• Use all enabled extensions

The following table outlines how user permissions determine what you can do with a extension:

Permissions Type	Permissions
Owners	• Edit extension name and description • Install the extension in a workspace or organization • Grant ownership and editing permissions to specific users and groups
Viewers	• View extension name and description • Access links to open in Slack, M365, and browser • Use the extensions

 Important

The Amazon Quick browser extension extension doesn't need any admin setup to install and use.

Extension access troubleshooting

If you encounter issues accessing or creating extensions:

- **Cannot create extensions:** Your IAM user who is a Amazon Quick administrator must first configure extension access for the extension type you want to use. Other administrators will not be able to create/edit/delete extension access.
- **Cannot find my extensions:** Check that extension access has been configured by your administrator and that you have the appropriate permissions to view extensions.
- **Extensions appear but cannot edit:** You have view-only access. Ask your administrator or extension owner to share edit permissions with you.
- **No extensions visible:** Contact your administrator to set up extension access and create default extensions for your organization.

Amazon Quick browser extension

Important

The Amazon Quick browser extension extension doesn't need any admin setup to install and use.

The Amazon Quick browser extension integrates AI-powered assistance directly into your web browsing experience. Users can access Quick capabilities within their browser to enhance productivity, streamline research, and get instant help without switching between applications.

Key capabilities include:

- Summarize web pages for quick understanding.
- Ask questions about multiple web pages and access organizational knowledge from spaces.
- Use actions from [connectors](#) configured in Quick.
- Analyze files uploaded to Quick during browsing sessions.

The AI-powered assistance enables you to maximize reading productivity, streamline research and analysis of complex information, and get instant help when creating content without context switching during web browsing.

Topics

- [Supported browsers](#)
- [Browser extension benefits](#)
- [Browser extension use cases](#)
- [Install and configure the browser extension](#)
- [Deploy the browser extension in your organization](#)
- [Browser extension usage guidelines](#)
- [Browser extension known limitations](#)
- [Using Amazon Quick Flows in the browser extension](#)

Supported browsers

The Quick browser extension is compatible with the following browsers:

- Google Chrome
- Mozilla Firefox
- Microsoft Edge

Browser extension benefits

The Quick browser extension provides powerful AI-driven capabilities that transform how you interact with web content. These key benefits help you work more efficiently and make better use of information you encounter while browsing, bringing enterprise knowledge directly into your web research workflows without requiring application switching.

Embedded productivity features

- Summarize a snapshot of any web page for quick understanding and analysis.
- Ask questions about one or more web pages using your organizational knowledge.
- Analyze both files you upload to Quick and specific spaces it can access.
- Execute actions from configured connectors without leaving your browser.

Browser extension use cases

The following are common use cases that help you make the best use of your Quick browser extension:

Summarize web pages

Use this procedure to quickly understand web page content without reading the entire page.

1. Open the Quick browser extension.
2. Log in and navigate to the web page you want to summarize.
3. From the chat interface, select **Summarize** for a summary of a snapshot of that web page.

Your conversation now contains a snapshot of this web page. You can continue to chat about the web page and ask follow-up questions.

Add individual pages to your conversation scope

Use this procedure to include multiple web pages in your conversation context for comprehensive analysis.

1. Navigate to the web page you want to add to your conversation.
2. Select the **+** icon to add the page to your scope.
3. Navigate from tab to tab and select the **+** icon for each page you want to chat about.
4. You can now chat about all the pages you've added to your scope.

This allows you to analyze and compare information across multiple web pages in a single conversation.

Add files to a conversation

Upload files to enhance your conversation with document analysis capabilities.

1. Select the **paperclip** icon.
2. Select the files to add to the conversation.
3. Select **Add to chat**.

The files are now available to use in the Quick conversation.

Add webpages as context to a conversation

Include multiple browser tabs in your conversation for comprehensive web content analysis.

1. Select **No tabs selected**.
2. From **Add tabs**, choose the tabs you want to include in your Quick conversation.
3. Select **Confirm**.

The tabs are now available in your Quick conversation.

Reset conversation context

Start fresh when you need to clear your current conversation context.

1. Select the **+ icon inside the bubble** to start a new conversation.
2. Alternatively, you can reset your conversation by selecting the **X** icon to end the current conversation and then open the Quick extension to start a new conversation.

This clears all previous context and allows you to begin with a clean slate.

Use actions in chat

Execute external actions directly from your browser extension conversation.

1. Ask the chat to perform a specific action.
2. Follow the prompts in the chat and on-screen to complete the action.

This enables you to perform tasks in external applications without leaving your browser.

Install and configure the browser extension

As a user, you can install and configure the browser extensions with the following steps.

- Mozilla based browsers — [Mozilla Firefox Add-ons web store](#)
- Chromium based browsers (including Microsoft Edge) — [Chrome Web Store](#)

Once you have successfully logged on, you can use your Quick browser extension.

Note

You can also **pin** your Quick browser extension to have it readily accessible while using your browser. Instructions for this are specific to your browser of choice. The following third-party information about pinning extensions might be helpful.

- Google Chrome — <https://www.howtogeek.com/683099/how-to-pin-and-unpin-extensions-from-the-chrome-toolbar>
- Mozilla Firefox — https://support.mozilla.org/en-US/kb/extensions-button#w_manage-pinned-extensions
- Microsoft Edge — <https://www.microsoft.com/en-us/edge/features/pin-to-taskbar>

Deploy the browser extension in your organization

As an author, after enabling the browser extension in your Quick application, you can deploy it across your organization using enterprise management tools and policies.

For Google Chrome Enterprise environments, you can use organizational unit policies to manage extension deployment. The Chrome Web Store for Enterprise provides additional deployment options specifically designed for organizational use. Mozilla Firefox Enterprise offers similar capabilities through policy templates and enterprise distribution methods that allow for automated extension deployment across the organization. Microsoft Edge Enterprise provides policy settings for managing extensions through mobile device management (MDM) software.

Policy settings from browser vendors: [Firefox](#), [Chrome](#), and [Edge](#).

Deploy to Google Chrome Enterprise

Follow these steps to deploy the browser extension across your Chrome Enterprise environment.

1. Access your Google Admin console and navigate to the Chrome management section.
2. Configure extension policies using organizational units to target specific user groups or departments.
3. Use the [Chrome extension management guide](#) to set up automated installation policies.
4. Reference the [Chrome Web Store Enterprise documentation](#) for advanced deployment configurations.

The extension will now be automatically installed for all users in your Chrome Enterprise environment.

Deploy to Mozilla Firefox Enterprise

Use these steps to deploy the extension across your Firefox Enterprise environment.

1. Configure Firefox deployment policies using the enterprise policy framework.
2. Follow the [Firefox extension deployment guide](#) for step-by-step instructions.
3. Use the [enterprise distribution documentation](#) to configure organization-wide deployment.
4. Apply extension settings using [Mozilla policy templates](#) for centralized configuration management.

The extension is now deployed across your Firefox Enterprise environment with centralized management.

Deploy to Microsoft Edge Enterprise

Deploy the extension to Microsoft Edge Enterprise using these organizational deployment steps.

1. Install the browser extension for all users using the software deployment processes of your organization.
2. Configure extension settings using [Edge policy settings](#) through mobile device management (MDM) software.
3. Follow the [Microsoft Edge extension management guide](#) for detailed deployment instructions.

Your Microsoft Edge Enterprise deployment is complete with centralized policy management.

Browser extension usage guidelines

As a user, you are responsible for keeping company information safe. The following guidelines help you use Quick browser extensions securely while maintaining data privacy and compliance.

Using the browser extension

When using the Quick browser extension, users can review past conversations and attachments to maintain context across browsing sessions.

Conversation retention and data privacy

Amazon Quick maintains strict data privacy through automated retention policies:

- Amazon Quick automatically deletes conversations and related files after 30 days of inactivity.
- Browser extension conversation retention follows your company's specific history rules.
- Browser extension retention periods may exceed Amazon Quick's standard 30-day limit based on organizational policies.

Browser extension known limitations

The following are known limitations of the Amazon Quick browser extension:

- The browser extension supports up to a total of 20 web pages and uploaded files at a time. For more information, see [Upload files and chat](#).

Using Amazon Quick Flows in the browser extension

Amazon Quick Flows in the browser extension enables you to run pre-defined workflows directly within your web browsing context. You can invoke flows that you've created or that have been shared with you, pass web page content as input, and execute multi-step workflows without leaving your browser.

When you use flows in the browser extension, you can:

- Invoke flows explicitly from the browser extension interface.
- Pass web pages as context to your flows.
- View flow execution history in your browser extension.

Prerequisites

Before using flows in the browser extension, ensure you have:

- Appropriate permissions to execute flows in your Amazon Quick account.

Invoking flows

You can invoke flows from the browser extension to run workflows on the web pages you're viewing.

To invoke a flow in the browser extension:

1. Open the Amazon Quick browser extension.
2. In the chat interface, choose the **Flows** icon.
3. Select a flow from the list or use the search bar to search through your flows.
4. The flow begins execution in the browser extension chat.
5. When prompted, select the tab you want to include as context and choose **Confirm**.

When a flow is active, the browser extension displays a progress tracker showing the current step. The dropdown menu on the progress tracker shows all steps in the flow.

Flow execution process

Once a flow is invoked in your browser extension, you interact with the flow through natural language chat as follows:

1. You're greeted with a summary of its capabilities.
2. You're prompted for required inputs through chat messages.
3. The extension displays outputs and results in the chat interface.
4. The extension shows a progress tracker indicating completed and remaining steps.

To view flow capabilities, see the [flows documentation](#).

Flow execution history

All flow executions in the browser extension are recorded in your conversation history. You can view past executions and resume incomplete flows. To view flow execution history:

1. In the browser extension, go to menu and choose the **Conversation history** icon.
2. Select a conversation that includes a flow execution.
3. The conversation displays the flow execution details, including all inputs and outputs.

If you closed the browser extension while a flow was running, you can resume the flow by selecting it from your history.

Stopping flow execution

You can stop a flow at any time during execution. To stop a running flow:

1. While the flow is running, choose the **End** button in the progress tracker.
2. The flow stops at the current step and you can continue your conversation with your My Assistant or Chat agent.

Amazon Quick Microsoft Outlook extension

The Amazon Quick extension for Outlook helps you manage your entire inbox using natural language. You can summarize unread messages, organize your inbox, schedule meetings, and draft responses, all without leaving Outlook.

You can search for Amazon Quick in the Microsoft Outlook app store or visit the [Quick for Outlook Microsoft store page](#) to add the extension.

Key capabilities include:

- **Inbox management:** Summarize and prioritize your unread emails, search for specific emails or discussions using natural language queries, and organize your inbox by moving emails to folders, flagging messages for follow-up, and more.
- **Calendar and meeting management:** Summarize your calendar, find optimal meeting times with coworkers, and schedule meetings using natural language instructions.
- **Email summaries and replies:** Open the Quick side panel to ask questions about an email thread in focus or generate a contextual reply.
- **Enterprise knowledge integration:** Use Quick's knowledge sources to draft contextual responses to emails or to perform inbox and calendar tasks.
- **External actions:** Perform actions in third-party applications using your configured [connectors](#) or apps directly from Outlook.

⚠ Important

- The Amazon Quick Outlook extension uses generative AI to create and execute code within your Outlook application sandbox. AI can make mistakes and perform inaccurate actions within your Outlook mailbox. No email or content is read when the side panel is closed, and no data is sent to Amazon Quick unless you explicitly send a prompt.
- Amazon Quick does not use your user data for service improvement or for training its underlying large language models (LLMs).

ℹ Note

Some features, such as inbox prioritization, calendar management, and email organization, require Microsoft Graph API permissions that must be granted by your administrator. If you are unable to perform these tasks, contact your administrator to ensure the required Graph API permissions have been approved for your organization. For more information, see [Microsoft Outlook extension permissions](#) in the admin guide.

Amazon Quick Microsoft Outlook extension usage guidelines

As a user, you are responsible for keeping company information safe. The following guidance helps you use Amazon Quick apps securely while maintaining data privacy and compliance.

Conversation retention and accuracy

Each conversation is stored for 30 days. You can review and manage your conversation history by choosing the conversation history button in the Outlook extension.

Amazon Quick uses generative AI. You should review responses for accuracy.

Usage of the Amazon Quick extension for Microsoft Outlook is subject to the [AWS Responsible AI Policy](#).

Security considerations

When you use Amazon Quick with email content, remember that responses reflect individual user permissions and content may include information not accessible to all email recipients. It may generate responses that contain sensitive data that can't be publicly shared.

To protect your organization's data, carefully evaluate email usage and plan deployment with data privacy in mind. Amazon Quick maintains strict data privacy by not using customer data for service improvements, not using customer data to enhance language models, and not indexing Microsoft Outlook conversations into your company's Amazon Quick instance.

Sample prompts

The following prompts demonstrate common ways to use Quick within Outlook. You can adapt these to your specific needs.

- "Summarize all unread emails and highlight any that need urgent responses today."
- "Draft a follow-up email to [contact name] referencing the key points from this thread."
- "Schedule a 30-minute meeting with the contacts in this email thread for next week."
- "Find all emails discussing [project name] and create a summary of action items."
- "Using data from my Quick space, draft a status update email for the leadership team on our quarterly metrics."

Tip

For best results, be specific in your prompts. Include names, dates, and goals. You can also reference your Quick spaces, dashboards, and knowledge bases for personalized results.

Amazon Quick Microsoft Excel extension

The Amazon Quick extension for Excel helps you work with data faster by automating data cleaning, analysis, and visualization tasks while integrating with your enterprise data sources.

You can search for Amazon Quick in the Microsoft Excel app store or visit the [Quick for Excel Microsoft store page](#) to add the extension.

Key capabilities include:

- **Data import and preparation:** Pull data from Quick dashboards, structured datasets, spaces, and uploaded files directly into your spreadsheet. Identify inconsistencies and clean data automatically, and apply formatting and templates based on your data and objectives.

- **Spreadsheet debugging:** Ask Quick to explain calculations behind cells or tables, debug and fix formula errors, and highlight anomalies and inconsistencies in your data.
- **Data analysis and insights:** Automatically analyze patterns and trends, create pivot tables and charts with natural language prompts, and perform sensitivity analysis and scenario planning.
- **Spreadsheet automation:** Apply filters, transformations, and conditional formatting using prompts, and generate appropriate templates based on your content and objectives.
- **Enterprise knowledge integration:** Choose an agent or a space to respond from and incorporate relevant information from your organization's knowledge into your spreadsheets. Integrate data from Quick dashboards, spaces, and knowledge bases.
- **External actions:** Perform actions in third-party applications using your configured [connectors](#) or apps directly from Excel.

Important

- The Amazon Quick Excel extension uses generative AI to create and execute code within your Excel application sandbox. AI can make mistakes and perform inaccurate actions within your Excel workbook. No spreadsheet data is read when the side panel is closed, and no data is sent to Amazon Quick unless you explicitly send a prompt.
- Amazon Quick does not use your user data for service improvement or for training its underlying large language models (LLMs).

Amazon Quick Microsoft Excel extension usage guidelines

As a user, you are responsible for keeping company information safe. The following guidance helps you use Amazon Quick apps securely while maintaining data privacy and compliance.

Conversation retention and accuracy

Each conversation is stored for 30 days. You can review and manage your conversation history by choosing the conversation history button in the Excel extension.

Amazon Quick uses generative AI. You should review responses for accuracy.

Usage of the Amazon Quick extension for Microsoft Excel is subject to the [AWS Responsible AI Policy](#).

Security considerations

To protect your organization's data, carefully evaluate usage and plan deployment with data privacy in mind. Amazon Quick maintains strict data privacy by not using customer data for service improvements, not using customer data to enhance language models, and not indexing Microsoft Excel conversations into your company's Amazon Quick instance.

Sample prompts

The following prompts demonstrate common ways to use Quick within Excel. You can adapt these to your specific needs.

- "Import data from my Quick dashboard and identify which items are trending above or below target."
- "Create a pivot table showing totals by category and region."
- "Generate a chart comparing this quarter's performance to last quarter."
- "Explain the formula in cell D5 and suggest a more efficient approach."
- "Apply conditional formatting to highlight any values that exceed the threshold by more than 10%."

Tip

For best results, be specific in your prompts. Include names, dates, and goals. You can also reference your Quick spaces, dashboards, and knowledge bases for personalized results.

Amazon Quick Microsoft Word extension

The Amazon Quick extension for Word helps you create professional documents faster by working directly in your document to generate, edit, and review content using natural language instructions.

You can search for Amazon Quick in the Microsoft Word app store or visit the [Quick for Word Microsoft store page](#) to add the extension.

Key capabilities include:

- **Content generation and editing:** Generate formatted content including headings, tables, and styled text using plain language instructions. Go from a blank page to a review-ready draft without manually formatting. Quick works directly in the document itself, automatically applying proper formatting and structure.
- **Document review with track changes:** Ask Quick to review your document and suggest edits based on any criteria or Quick data of your choice. Suggestions appear in Word's built-in track changes and comments, so you stay in control of what gets accepted.
- **Style and formatting matching:** Quick automatically matches your document's existing style, structure, and formatting. Attach reference documents or templates to apply specific formatting standards.
- **Enterprise knowledge integration:** Include relevant information from your Quick knowledge sources directly into your document.
- **External actions:** Perform actions in third-party applications using your configured [connectors](#) or apps directly from Word.

Important

- The Amazon Quick Word extension uses generative AI to create and execute code within your Word application sandbox. AI can make mistakes and perform inaccurate actions within your Word document. No document content is read when the side panel is closed, and no data is sent to Amazon Quick unless you explicitly send a prompt.
- Amazon Quick does not use your user data for service improvement or for training its underlying large language models (LLMs).

Amazon Quick Microsoft Word extension usage guidelines

As a user, you are responsible for keeping company information safe. The following guidance helps you use Amazon Quick apps securely while maintaining data privacy and compliance.

Conversation retention and accuracy

Each conversation is stored for 30 days. You can review and manage your conversation history by choosing the conversation history button in the Word extension.

Amazon Quick uses generative AI. You should review responses for accuracy.

Usage of the Amazon Quick extension for Microsoft Word is subject to the [AWS Responsible AI Policy](#).

Security considerations

To protect your organization's data, carefully evaluate usage and plan deployment with data privacy in mind. Amazon Quick maintains strict data privacy by not using customer data for service improvements, not using customer data to enhance language models, and not indexing Microsoft Word conversations into your company's Amazon Quick instance.

Sample prompts

The following prompts demonstrate common ways to use Quick within Word. You can adapt these to your specific needs.

- "Create a project proposal including an executive summary, timeline, and budget section, pulling data from my Quick space."
- "Review this document with Track Changes enabled and suggest edits to improve clarity and tone."
- "Read the comments in this document and address each piece of feedback, making the requested changes directly."
- "Generate a one-page executive summary of this report highlighting the key findings."
- "Add a competitive comparison table to this document using data from my Quick dashboard."

Tip

For best results, be specific in your prompts. Include names, dates, and goals. You can attach reference documents or templates to guide formatting and style.

Amazon Quick Microsoft PowerPoint extension

The Amazon Quick extension for PowerPoint helps you create compelling, data-driven presentations by integrating your enterprise data, AI-powered design assistance, and intelligent content generation directly into PowerPoint.

You can search for Amazon Quick in the Microsoft PowerPoint app store or visit the [Quick for PowerPoint Microsoft store page](#) to add the extension.

Key capabilities include:

- **Presentation creation with enterprise data:** Generate complete presentations using data from Quick dashboards, spaces, and knowledge bases. Pull real-time metrics and insights directly from your organization's data sources, or create presentations from outlines, topics, or natural language descriptions.
- **Visual narratives from enterprise knowledge:** Use Quick's enterprise knowledge to create contextual, accurate presentations. Import and visualize data from your Quick spaces, dashboards, and knowledge bases in compelling slides.
- **Brand consistency and professional design:** Apply your organization's themes and templates to presentations, respect existing presentation formatting and primary slide layouts, and generate professionally formatted tables, layouts, and design elements.
- **Natural language editing:** Select slides or content and request modifications using conversational prompts. Add, delete, reorder, or reorganize slides effortlessly, and chat about your presentation content to get insights and suggestions.

Important

- The Amazon Quick PowerPoint extension uses generative AI to create and execute code within your PowerPoint application sandbox. AI can make mistakes and perform inaccurate actions within your PowerPoint slide deck. No presentation data is read when the side panel is closed, and no data is sent to Amazon Quick unless you explicitly send a prompt.
- Amazon Quick does not use your user data for service improvement or for training its underlying large language models (LLMs).

Amazon Quick Microsoft PowerPoint extension usage guidelines

As a user, you are responsible for keeping company information safe. The following guidance helps you use Amazon Quick apps securely while maintaining data privacy and compliance.

Conversation retention and accuracy

Each conversation is stored for 30 days. You can review and manage your conversation history by choosing the conversation history button in the PowerPoint extension.

Amazon Quick uses generative AI. You should review responses for accuracy.

Usage of the Amazon Quick extension for Microsoft PowerPoint is subject to the [AWS Responsible AI Policy](#).

Security considerations

To protect your organization's data, carefully evaluate usage and plan deployment with data privacy in mind. Amazon Quick maintains strict data privacy by not using customer data for service improvements, not using customer data to enhance language models, and not indexing Microsoft PowerPoint conversations into your company's Amazon Quick instance.

Sample prompts

The following prompts demonstrate common ways to use Quick within PowerPoint. You can adapt these to your specific needs.

- "Create a 10-slide presentation on [topic] including key metrics and recommendations, pulling data from my Quick dashboard."
- "Add a slide with a competitive comparison showing our advantages over the top 3 competitors."
- "Make this deck more executive-friendly by reducing text and emphasizing key numbers."
- "Reorganize this presentation to lead with results instead of methodology."
- "Create a chart slide comparing actual vs. target performance by department using data from my Quick space."

Tip

For best results, be specific in your prompts. Include names, dates, and goals. You can also reference your Quick spaces, dashboards, and knowledge bases for personalized results.

Amazon Quick Slack extension

Note

Before you deploy a Slack extension as an author, your Amazon Quick admin must [configure Amazon Quick access to Slack](#).

The Amazon Quick Slack extension brings AI-powered assistance directly into team conversations, enabling access to organizational knowledge and document analysis without leaving Slack.

Key capabilities include:

- Extension access through Slack's AI assistant side panel for continuous availability.
- Engagement with Quick through **@Amazon Quick** mentions in channels for contextual team assistance. Contextual responses with relevant content from the Slack channel are not supported at this time.
- Direct interaction through private messages for individual productivity.
- External actions in third-party applications from Slack using [connectors](#) configured by your organization.
- Unlimited usage for enhanced team collaboration and knowledge sharing.
- General knowledge capabilities.
- Company-specific information search when enabled.
- Source-attributed responses.
- File upload and processing of up to 20 files simultaneously.
- Contextual summaries and insights generation.

Topics

- [Amazon Quick Slack extension author guide](#)
- [Amazon Quick Slack extension user guide](#)

Amazon Quick Slack extension author guide

As an Amazon Quick author, you can deploy Quick Slack extensions after your admin establishes the foundational connection to your organization's Slack workspace. Your capabilities depend on the permission level granted by your administrator.

With **limited permissions** (view, share, delete only), you can manage basic extension operations through the landing page after admin completes all setup. With **full permissions** (deploy, view, share, delete, edit), you can complete the Slack OAuth application deployment, rename extensions, and access all editing features.

Author capabilities for Slack extensions:

- Deploy extensions to your organization's Slack workspace using OAuth (requires full permissions).
- Customize extension settings - names, descriptions (requires full permissions).
- Manage sharing and access permissions (available with limited or full permissions).

Note

Before you deploy a Slack extension as an author, your Quick admin must [configure Amazon Quick access to Slack](#).

Topics

- [Deploy a Slack extension](#)
- [Edit Slack extension](#)
- [Share Slack extension](#)
- [Delete Slack extension](#)

Deploy a Slack extension

Deploy a new Slack extension instance in the Quick console. This process establishes the foundation for connecting AI-powered assistance to your Slack workspace.

Note

This action requires full author permissions.

1. Sign in to the Amazon Quick console.
2. In the left navigation, under **CONNECTIONS**, select **Extensions**.
3. Select **Create extension**.
4. Select **Slack**. Then, select **Next**.
5. Configure the following fields:
 - **Name** - A name for your extension is pre-filled for you. You can edit this and enter a descriptive name for the Slack extension.

- **Description** (optional) - A description for your extension is pre-filled for you. You can edit this and enter a new description to provide additional context about this extension configuration.
 - **Installation** type - Your Slack extension supports OAuth installation by default.
6. Select **Next** to save your configuration.
 7. From the **Extension** summary page, navigate to the extension you just configured.
 8. Then, from the **Actions** menu, navigate to the extension you just configured.
 9. Select **Install**. Then, from the **Complete installation for Slack assistant** dialog box that opens, select **Install**.

A success message will open up on the top right of your screen.

10. From the success message, select **Install extension** to finish installing your extension.

 Note

You can also navigate to the extensions summary page and deploy your extension from the **Actions** menu.

11. The link will open to a Slack login page where after you login (as admin) you will need to find and be asked to install the Amazon Quick app within your Slack workspace.

 Note

You may see a "This app is not approved by Slack banner." This message can be ignored.

12. Choose **Allow** to install your Amazon Quick App for Slack.
13. Once the installation is complete, you will see the page confirming that the **Congratulations! Your Slack App has been successfully installed.**
14. Choose **Open the Amazon Quick App in Slack.**
15. This will open your Slack workspace where all users will be required to sign-in.

Your Slack extension has now been successfully deployed and is available for users.

Edit Slack extension

Note

This action requires full author permissions.

As an author, you can edit the extensions you deploy to your users. Modify extension settings to update names, descriptions, or configuration options. Changes take effect immediately and apply to all users with access to the extension.

1. Sign in to the Amazon Quick console.
2. In the left navigation, under **CONNECTIONS**, select **Extensions**.
3. Select the three dot menu icon for the Slack extension you need to edit.
4. Select **Edit**.
5. Edit the configuration as required and select **Save** to confirm the changes.

Share Slack extension

Share ownership and management permissions with specific users and groups, enabling multiple users to manage extensions and assist with installation. You can assign different permission levels and manage access as needed.

1. Sign in to the Amazon Quick console.
2. In the left navigation, under **CONNECTIONS**, select **Extensions**.
3. Select the three dot menu icon for the Slack extension you need to share.
4. Select **Share**.
5. Enter the users and groups you would like to share the extension with.
6. Select **Share** to send the access email to each group and user.
7. From the drop-down next to each name, you can edit their access levels (**Viewer** or **Owner**).
8. **Optional:** You could select **Remove access** to delete the access for the selected group or user.

Delete Slack extension

As an author, you can delete the extensions you deploy to your users. Permanently remove a extension from your Quick console and revoke access for all users. This action cannot be undone and requires confirmation.

1. Sign in to Amazon Quick console.
2. In the left navigation, under **CONNECTIONS**, select **Extensions**.
3. Select the three dot menu icon for the Slack extension you need to delete.
4. Select **Delete**.
5. Enter the word, "confirm", and select **DELETE**.

Amazon Quick Slack extension user guide

As a user, you can access Amazon Quick directly within Slack to get AI-powered assistance without leaving your workspace. The extension integrates seamlessly into your Slack environment, providing instant access to knowledge bases and intelligent responses.

Topics

- [Add a Slack Extension to your Slack](#)
- [Access Quick in your Slack side panel at any time](#)
- [Slack extension use cases](#)
- [Amazon Quick Slack extension usage guidelines](#)
- [Slack extension known limitations](#)

Add a Slack Extension to your Slack

Install the Amazon Quick app in your Slack workspace to enable AI-powered assistance. This one-time setup makes Quick accessible throughout your Slack environment.

Add a Slack Extension to your Slack

1. Open and login to the Slack workspace for your company.
2. Choose **More** in the left navigation, then select **Automations**.
3. Choose **Apps**.

4. Choose **+ Add apps**.
5. Search for **Quick** and choose Quick.
6. Select your Slack profile picture in the bottom left of Slack.
7. Navigate to **Preferences > Navigation > AI Apps and Agents** > select **Show app agents** > and select **Quick**.
8. Quick is now available to access anywhere in Slack through an icon in the top right of your Slack window.

Note

For more information, see [Understand AI apps in Slack](#) in the Slack help center.

Access Quick in your Slack side panel at any time

Once you've added Quick to your Slack workspace, you can access it conveniently through the side panel for quick assistance without interrupting your workflow. The side panel provides persistent access to Quick while you work in Slack channels and direct messages.

Access Quick in your Slack side panel at any time

1. Click on your Slack profile picture in the bottom left of Slack.
2. Navigate to **Preferences > Navigation > AI Apps and Agents** and select **Amazon Quick**.
3. Quick will now be available to access anywhere in Slack through an icon in the top right of your Slack window.

Slack extension use cases

With Quick integrated into Slack, you can search and access your organization's knowledge bases directly from your conversations. This allows you to find relevant information, get answers to questions, and enhance your communications without leaving Slack.

Search your organization's knowledge wherever you are

1. Open the Quick side panel from the top right of your Slack window or start a DM with Quick under **Apps**.

2. Ask Quick a question about its knowledge from your spaces and connected knowledge bases.
3. To personalize your conversation using your data, use the gear icon (visible after sending your first message to Quick) within your conversation to select an agent or a space to respond from.

Get Quick's in-context help in Slack conversation threads

1. Navigate to a Slack conversation thread of your choice.
2. Type **@Amazon Quick** and ask any question about the messages in that thread or its knowledge from your spaces and Quick knowledge bases.
3. If Quick is not already added to the channel, you will be prompted to add it before receiving a response.

Perform actions in external applications

1. Start a direct message (DM) with Quick.
2. Ask it to perform an action of your choice in an external application using [connectors](#) configured by your organization.

Upload and analyze images

1. In any Slack channel or direct message, upload an image by dragging and dropping it or using the attachment button.
2. Mention **@Amazon Quick** and ask questions about the uploaded image, such as requesting analysis, descriptions, or insights.
3. Quick will analyze the image and provide relevant information based on your request.

Amazon Quick Slack extension usage guidelines

As a user, you are responsible for keeping company information safe. The following guidance helps you use Amazon Quick apps securely while maintaining data privacy and compliance.

When using the Slack extension

To start fresh, begin a new conversation by selecting **New chat** in direct messages. Closing the side panel ends your current conversation. You can access conversation history through Slack message history or the Amazon Quick chat instance, which includes conversations from all channels. You can

view complete conversation details, including attachment names, and access the same knowledge base available in your Amazon Quick chat instance.

Note

Users can access conversation history in Slack. Deleting the conversation in Slack will not delete from conversation history in Amazon Quick. However, users need to visit the Amazon Quick chat instance to manage their conversations.

Conversation retention policy and accuracy

Amazon Quick automatically deletes conversations after 30 days of inactivity. Slack conversation retention follows your company's specific history rules, and Slack retention periods may exceed Amazon Quick's 30-day limit.

Users can access conversation history in Slack through the AI App History tab. Deleting the conversation in Slack will not delete from conversation history in Amazon Quick. However, users need to visit the Amazon Quick chat instance to manage their conversations.

Security considerations

When you use Amazon Quick in public channels, remember that responses reflect individual user permissions and content may include information not accessible to all channel members. It may generate responses that contain sensitive data that can't be publicly shared.

To protect your organization's data, carefully evaluate public channel usage, consider the security implications, and plan deployment with data privacy in mind. Amazon Quick maintains strict data privacy by not using customer data for service improvements, not using customer data to enhance language models, and not indexing Slack conversations into your company's Amazon Quick instance.

Slack extension known limitations

The following are known limitations of the Amazon Quick Slack extension:

- The Slack extension doesn't have access to all your Slack conversations. It only has access to the messages in Slack conversation threads where it is mentioned.
- The Slack extension doesn't have access to previous messages when mentioned in a Slack channel or in group message.

- The Slack extension isn't accessible in group direct messages or in direct messages with teammates unless it is explicitly added to the conversation as a member.
- You can filter Amazon Quick agents and spaces only in direct messages with the Slack extension. The Slack extension will use the default assistant and all spaces available to you when invoked in public such as in channels or group messages.
- The Slack extension can't render visuals for prompts about your structured data.
- Web search isn't supported in the Slack extension.
- The Slack extension can't be configured to automatically reply to messages in channels or posts.
- The Slack extension will not respond when invoked by Slack's workflow bot as it is not able to authenticate the user who made the query in these situations.
- The Slack extension isn't available in huddles.
- [Actions](#) (using your action integrations) can only be performed in direct messages with the Slack extension.
- Actions that require file uploads as inputs are not supported by the Slack extension.
- [Flows](#) are not supported in the Slack extension at this time.
- File upload limitations are the same within Amazon Quick chat. For more information, see [Upload files and chat](#).

Amazon Quick Microsoft Teams extension

Note

Before you deploy a Microsoft Teams extension as an author, your Quick admin must [configure Amazon Quick access to Microsoft Teams](#).

The Amazon Quick Microsoft Teams extension brings AI-powered assistance directly into team conversations, enabling access to organizational knowledge and document analysis without leaving Microsoft Teams.

Key capabilities include:

- Seamless integration within Microsoft Teams allowing you to access information while staying in your conversations.

- Collaborative assistance through **@Amazon Quick** mentions in thread discussions for thread-specific information.
- Action item extraction and summary generation from conversations.
- File upload and processing capabilities for enhanced productivity.

Topics

- [Amazon Quick Microsoft Teams extension author guide](#)
- [Amazon Quick Microsoft Teams extension user guide](#)

Amazon Quick Microsoft Teams extension author guide

As an Amazon Quick author, you can deploy Quick Microsoft Teams extensions after your admin establishes the foundational connection to your organization's Microsoft Teams workspace. Your capabilities depend on the permission level granted by your administrator.

With **limited permissions** (view, share, delete only), you can manage basic extension operations through the landing page after admin completes all setup. With **full permissions** (deploy, view, share, delete, edit), you can complete OAuth deployments, download manifests, rename extensions, and access all editing features.

Author capabilities for Microsoft Teams extensions include:

- Configure extension connections (requires admin privileges).
- Deploy extensions to your organization's Microsoft Teams workspace (requires full permissions).
- Manage sharing and access permissions (available with limited or full permissions).
- Customize extension settings - names, descriptions (requires full permissions).
- Download manifest files for advanced Microsoft Teams customization (requires full permissions).

Note

Before you deploy a Microsoft Teams extension as an author, your Quick admin must [configure Amazon Quick access to Microsoft Teams](#).

Topics

- [Deploy Microsoft Teams extension](#)
- [Edit Microsoft Teams extension](#)
- [Share Microsoft Teams extension](#)
- [Delete Microsoft Teams extension](#)

Deploy Microsoft Teams extension

Deploy a new Microsoft Teams extension instance in the Amazon Quick console. This process establishes the foundation for connecting AI-powered assistance to your Microsoft Teams workspace.

Note

This action requires full author permissions.

Use this procedure to create and configure a new Microsoft Teams extension for your organization.

1. Sign in to the Amazon Quick console.
2. In the left navigation, under **CONNECTIONS**, select **Extensions**.
3. Select **Create extension**.
4. Select **Microsoft Teams**. Then, select **Next**.
5. Configure the following fields:
 - **Name** - A name for your extension is pre-filled for you. You can edit this and enter a descriptive name for the Microsoft Teams extension.
 - **Description** (optional) - A description for your extension is pre-filled for you. You can edit this and enter a new description to provide additional context about this extension configuration.
 - **Installation** type - Your Microsoft Teams extension supports OAuth installation by default.
6. Select **Next** to save your configuration.
7. From the **Extension** summary page, navigate to the extension you just configured.
8. Then, from the **Actions** menu, navigate to the extension you just configured.
9. Select **Install**. Then, from the **Complete installation for Teams assistant** dialog box that opens, select **Install**.

A success message will open up on the top right of your screen.

10. From the success message, select **Install extension** to finish installing your extension.

 Note

You can also navigate to the extensions summary page and deploy your extension from the **Actions** menu.

11. Open the link and login as Global Admin or someone with administrative permissions can add the Amazon Quick App to the Microsoft Teams admin center for your organization.
12. Choose **Teams apps** in the left navigation.
13. Choose **Amazon Quick** from the list of available apps.
14. Review and grant admin consent by choosing the **Permissions** tab and reviewing the permissions and choose **Grant admin consent**.

 Note

If permissions are already granted, proceed to the end of the procedure, there is no further action required.

15. Authenticate and choose **Accept** for Amazon Quick app.
16. Confirm that an app titled Amazon Quick **Permissions** tab now says **Admin consent granted for all required permissions**.
17. All users assigned to the app from the Teams admin center can now find the app in the **Built for your org** section of the **Apps** page of their Teams app.

Your Microsoft Teams extension has now been successfully deployed and is available for users.

Edit Microsoft Teams extension

 Note

This action requires full author permissions.

As an author, you can edit the extensions you deploy to your users. Modify extension settings to update names, descriptions, or configuration options. Changes take effect immediately and apply to all users with access to the extension.

Use this procedure to modify settings and configuration for an existing Microsoft Teams extension.

1. Sign in to the Amazon Quick console.
2. In the left navigation, under **CONNECTIONS**, select **Extensions**.
3. Select the three dot menu icon for the Microsoft Teams extension you need to edit.
4. Select **Edit**.
5. Edit the configuration as required and select **Save** to confirm the changes.

Your changes are now applied and will be reflected in the extension configuration for all users.

Share Microsoft Teams extension

Share ownership and management permissions with specific users and groups, enabling multiple users to manage extensions and assist with installation. You can assign different permission levels and manage access as needed.

Use this procedure to share your Microsoft Teams extension with other users and manage their access permissions.

1. Sign in to the Amazon Quick console.
2. In the left navigation, under **CONNECTIONS**, select **Extensions**.
3. Select the three dot menu icon for the Microsoft Teams extension you need to share.
4. Select **Share**.
5. Enter the users and groups you would like to share the extension with.
6. Select **Share** to send the access email to each group and user.
7. From the drop-down next to each name, you can edit their access levels (**Viewer** or **Owner**).
8. **Optional:** You could select **Remove access** to delete the access for the selected group or user.

The specified users and groups now have access to your Microsoft Teams extension with the permissions you assigned.

Delete Microsoft Teams extension

As an author, you can delete the extensions you deploy to your users. Permanently remove an extension from your Quick console and revoke access for all users. This action cannot be undone and requires confirmation.

Use this procedure to permanently remove a Microsoft Teams extension from your organization.

1. Sign in to Amazon Quick console.
2. In the left navigation, under **CONNECTIONS**, select **Extensions**.
3. Select the three dot menu icon for the Microsoft Teams extension you need to delete.
4. Select **Delete**.
5. Enter the word, "confirm", and select **DELETE**.

The Microsoft Teams extension has been permanently removed and is no longer accessible to users in your organization.

Amazon Quick Microsoft Teams extension user guide

As a user, you can access Quick directly within Microsoft Teams to get AI-powered assistance without leaving your workspace. The extension integrates seamlessly into your Microsoft Teams environment, providing instant access to knowledge bases and intelligent responses.

Topics

- [Add a Microsoft Teams Extension to your Microsoft Teams](#)
- [Microsoft Teams extension use cases](#)
- [Amazon Quick Microsoft Teams extension usage guidelines](#)
- [Microsoft Teams extension known limitations](#)

Add a Microsoft Teams Extension to your Microsoft Teams

Install the Quick app in your Microsoft Teams workspace to enable AI-powered assistance. This one-time setup makes Quick accessible throughout your Microsoft Teams environment.

Add a Microsoft Teams Extension to your Microsoft Teams

1. Open and login to the Microsoft Teams workspace for your company.

2. Choose **Apps** in the left navigation and search for **Quick**.
3. Choose **Quick**.
4. On the top bar, Quick icon is now available. You can select it to access and chat with Quick.

Note

For more information, see [Chat with a bot in Microsoft Teams](#) from Microsoft support.

Microsoft Teams extension use cases

With Amazon Quick integrated into Microsoft Teams, you can search and access your organization's knowledge bases directly from your conversations. This allows you to find relevant information, get answers to questions, and enhance your communications without leaving Team.

Search your organization's knowledge without leaving Teams

1. Start a chat with Quick after adding it from the **Apps** section.
2. Ask Quick a question about its knowledge from your spaces and connected knowledge bases.
3. To personalize your conversation using your data, use the gear icon (visible after sending your first message to Quick) within your conversation to select an agent or a space to respond from.

Get Quick's in-context help in Teams posts

1. Navigate to a Teams post of your choice.
2. Type **@Amazon Quick** and ask any question about the messages in that thread or its knowledge from your spaces and your organization's knowledge bases.
3. If Quick is not already added to the channel, you will be prompted to add it before receiving a response.

Perform actions in external applications

1. Start a direct message (DM) with Quick.
2. Ask it to perform an action of your choice in an external application using [connectors](#) configured by your organization.

Amazon Quick Microsoft Teams extension usage guidelines

As a user, you are responsible for keeping company information safe. The following guidance helps you use Amazon Quick apps securely while maintaining data privacy and compliance.

Conversation retention policy and accuracy

Amazon Quick automatically deletes conversations after 30 days of inactivity. Teams conversation retention follows your company's specific history rules. Teams retention periods may exceed Amazon Quick's 30-day limit. To start fresh, begin a new conversation by using the `/new_conv` command

Users can review past conversations in Teams or all conversations from all channels (Teams, browser extensions, etc.) in your Amazon Quick chat instance.

Note

Deleting conversations or messages in Teams does not remove them from Amazon Quick. To manage your conversation history, use the Amazon Quick chat instance.

Security considerations

When Amazon Quick is invoked by a user in a public Teams channel, it generates responses based on the invoking user's permissions, which may include content that other channel members aren't authorized to access. To prevent unintended exposure of sensitive information, carefully evaluate the use of Amazon Quick in public channels.

File upload behavior in Amazon Quick Teams Extension:

- Responses will be limited to the uploaded file content.
- General knowledge access depends on admin settings.
- Start a new chat to access company knowledge again.

Data privacy and conversation behavior

Amazon Quick does not use customer data for service improvement or for improving its underlying large language models (LLMs). Also, none of the data you include in your Teams conversations will be indexed into your company's Amazon Quick instance.

The Amazon Quick Teams app will have access to the same knowledge available in the corresponding Amazon Quick web experience.

Closing the Amazon Quick bot for Teams side panel will end the current conversation. Users can review past conversations in Teams or all conversations from all channels (Teams, browser extensions, etc.) in your Amazon Quick web experience. You can access all the history of previous conversations including, the names of the attachments in those conversations.

Microsoft Teams extension known limitations

The following are known limitations of the Amazon Quick Microsoft Teams extension:

- The Microsoft Teams extension doesn't have access to all your Microsoft Teams conversations. It only has access to the messages in Microsoft Teams posts where it is mentioned.
- The Microsoft Teams extension isn't accessible in group direct messages or in direct messages with other teammates.
- You can filter Amazon Quick agents and spaces only in direct messages with the Microsoft Teams extension. The Microsoft Teams extension will use the default assistant and all spaces available to you when invoked in public, such as in channels and Microsoft Teams posts.
- The Microsoft Teams extension can't render visuals for prompts about your structured data.
- Web search is not supported in the Microsoft Teams extension.
- The Microsoft Teams extension can't be configured to automatically reply to messages in channels or posts.
- The Microsoft Teams extension isn't available for your Microsoft Teams meetings.
- You cannot review your conversation history in Microsoft Teams. To review and manage your conversation history, go to your Amazon Quick chat instance.
- [Actions](#) (using your action integrations) can only be performed in direct messages with the Microsoft Teams extension.
- Actions that require file uploads as inputs are not supported by the Microsoft Teams extension at this time.
- [Flows](#) are not supported in the Microsoft Teams extension at this time.
- File upload limitations are the same as those in the web experience. For more information see [Upload files and chat](#).

Amazon Quick on desktop

Amazon Quick on desktop is an agentic AI teammate that connects to your local files and the tools you use every day. It builds an understanding of your work that deepens over time, so it provides relevant, personalized assistance. Agents take on work you delegate, follow through independently, and bring you the results.

For the plans that include Amazon Quick on desktop and pricing details, see [Amazon Quick pricing](#).

Topics

- [What is Amazon Quick on desktop?](#)
- [Setting up Amazon Quick on desktop for enterprise deployments](#)
- [Getting started](#)
- [Understanding the desktop interface](#)
- [Working with chat](#)
- [Activity feed](#)
- [Library](#)
- [Apps in Amazon Quick](#)
- [Spaces](#)
- [Customize](#)
- [Agents](#)
- [Skills](#)
- [Connectors](#)
- [Knowledge](#)
- [Mission Control](#)
- [My context](#)
- [Memories](#)
- [Knowledge graph](#)
- [Settings](#)
- [System tools](#)

- [Security, privacy, and architecture](#)
- [Troubleshooting](#)

What is Amazon Quick on desktop?

Amazon Quick on desktop is a native application that extends Amazon Quick to your machine. It gives Quick direct, persistent access to your files, your connected services, and the context of how your work fits together.

What you can do with Quick on desktop

Quick on desktop works with the files already on your machine. You grant access to the folders you choose, and Quick reads, writes, searches, and indexes them directly.

Amazon Quick runs in the cloud. Your conversations, memory, knowledge graph, agent-produced files, and the index of content you make searchable are kept for you in your Amazon Quick account, so they are available across the web, mobile, and desktop. What the desktop application adds is access to your machine: the files in the folders you grant, browser automation in your own browser, local code execution, and MCP servers and coding agents you configure locally. The desktop keeps only what it needs to work locally: your folder permissions, local connector configuration, and sign-in credentials.

Quick connects to the services you use every day: email, calendar, chat, customer relationship management (CRM) tools, and more. Once connected, Quick pulls information from these services, takes actions in them, and combines data across them in a single conversation.

As you work with Quick, it builds a knowledge graph of the people, projects, and decisions relevant to you, and it retains your preferences and workflows in long-term memory. Together, these let Quick understand your work context and provide personalized responses that improve over time. You review and manage both from My context.

You create agents with their own instructions and knowledge, and pair them with schedules that run at intervals you define. Agents prepare for your meetings, summarize what changed in your channels, generate reports, and deliver results through an activity feed.

Quick on desktop is extensible. You connect MCP servers to add custom tools and integrations, delegate development tasks to coding agents, use browser automation to interact with web applications, and create skills that teach Quick new workflows.

Next steps

To get started, see [Getting started](#). If your organization uses an Enterprise account, see [Setting up Amazon Quick on desktop for enterprise deployments](#) first.

Setting up Amazon Quick on desktop for enterprise deployments

Applies to: Enterprise Edition

Intended audience: System administrators

To use Amazon Quick on desktop for enterprise deployments, administrators must configure enterprise single sign-on (SSO) so that users in the organization can sign in with their corporate credentials. This setup connects your organization's OpenID Connect (OIDC) compatible identity provider (IdP) to Amazon Quick.

Note

If you are using a Free or Plus account, this section does not apply to you. Continue to [Getting started](#).

How enterprise sign-in works

The Amazon Quick desktop application uses the OIDC protocol to authenticate users. When you choose **Continue with SSO**, a browser window opens to Amazon Quick. Amazon Quick identifies your configured identity provider and redirects to your IdP's authorization endpoint. The application then exchanges the resulting authorization code for tokens using Proof Key for Code Exchange (PKCE).

To identify the configured identity provider, your browser must have an active Amazon Quick web session. If no session exists, the browser prompts you to sign in to Amazon Quick before the application redirects to your IdP.

Amazon Quick validates the token and maps the user to an identity in your account. The email address in your IdP must exactly match the email address of the user in Amazon Quick.

 Note

If you sign in for the first time, the browser checks for an active Amazon Quick web session. If no session exists, the browser opens the Amazon Quick sign-in page instead of redirecting to your IdP. Sign in to Amazon Quick in the browser to continue. The desktop application then completes sign-in through your configured identity provider.

Supported token signing algorithms

Amazon Quick verifies the signature on each ID token against the public keys published at your identity provider's JWKS URI. Your IdP must sign ID tokens with one of the following asymmetric algorithms.

Key type	Supported alg values
RSA	RS256, RS384, RS512, PS256, PS384, PS512
ECDSA	ES256, ES384, ES512

 Symmetric signing is not supported

Amazon Quick does not accept ID tokens signed with HMAC (HS256, HS384, or HS512). These algorithms sign the token with the client secret rather than with a key that can be published in a JWKS, so the signature cannot be verified against your IdP's JWKS URI. Unsecured tokens (an alg value of none) are also rejected. If your IdP is configured to use one of these, sign-in fails with a token validation error.

You do not select the algorithm in Amazon Quick. Your IdP chooses it when it issues the token, and Amazon Quick accepts any of the algorithms in the preceding table. Most identity providers use RS256 by default.

Prerequisites

Before you begin, verify that you have the following:

- An AWS account with an active Amazon Quick subscription. The Amazon Quick account's home region (identity region) must be in a supported AWS Region. For a list of supported Regions, see [Supported AWS Regions for Amazon Quick](#). All identity types are supported, including IAM Identity Center, IAM federation, and native Amazon Quick (username/password) users.
- Administrator access to your Amazon Quick account.
- Access to your IdP with permissions to create OIDC application registrations.
- In restricted network environments, the ability to reach the required Amazon Quick and identity provider domains. For the list of domains to add to your allow list, see [Network access and required domains](#).

Important

Amazon Quick on desktop is available for Enterprise accounts in AWS Regions that support the full Amazon Quick feature set. Regions that support Amazon Quick Sigh capabilities only do not include desktop. For the full list, see [Supported AWS Regions for Amazon Quick](#).

Setup process

Setting up enterprise sign-in involves the same four steps regardless of which identity provider you use:

1. **Create an OIDC application in your identity provider.** Register a public OIDC client and record its Client ID and OIDC endpoints. The steps and endpoint formats are specific to your identity provider.
2. **Add the extension access in the Amazon Quick administration console.** In the Amazon Quick administration console, add an extension access using the Client ID and OIDC endpoints from Step 1.
3. **Create the extension in the Amazon Quick console.** On the **Extensions** page in the Amazon Quick console, create the extension from the extension access you added in Step 2.

4. **Download, verify, and distribute the desktop application.** Download and install the application, choose **Continue with SSO** to confirm that authentication succeeds, and then direct your users to download and sign in.

The complete procedure for each of these steps is documented on the page for your identity provider. Choose your identity provider to get started:

- [Microsoft Entra ID](#) – Microsoft Entra ID
- [Google Workspace](#) – Google Workspace
- [Okta](#) – Okta
- [Ping Identity](#) – Ping Identity (PingFederate and PingOne)

If you encounter problems during setup or sign-in, see [Troubleshooting enterprise sign-in for Amazon Quick on desktop](#).

Managed deployment configuration

Administrators can configure and deploy Amazon Quick on desktop across an enterprise fleet using the following options.

Managed policies

Amazon Quick reads policies from vendor-neutral OS-managed locations. Any mobile device management (MDM) tool that can deliver a configuration profile (macOS) or registry policy (Windows) works without vendor-specific integration.

Policy	macOS location	Windows location	Effect
DisableSocialLogin	Preference domain com.aws.Q uickWork.mac	HKLM\SOFTWARE \Policies \Amazon\Quick (REG_DWORD)	Hides social sign-in, blocks it server-side, and hides "Sign up for free." This policy enforces enterprise SSO as the only authentication path.

Policy	macOS location	Windows location	Effect
DisableAutoUpdates	Preference domain com.aws.QuickWork.mac	HKLM\SOFTWARE\Policies\Amazon\Quick (REG_SZ)	Stops the application from updating itself, so that your organization controls when new versions reach the fleet. Accepts patch, minor, or major. The value is the smallest update size that the policy disables. patch disables all automatic updates, including patch, minor, and major updates. minor disables minor and major updates. major disables major updates only.

Amazon Quick reads policy values at application startup.

macOS

Deploy a configuration profile with preference domain com.aws.QuickWork.mac. The profile supports device scope (/Library/Managed Preferences/com.aws.QuickWork.mac.plist) or user scope.

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE plist PUBLIC "-//Apple//DTD PLIST 1.0//EN"
    "http://www.apple.com/DTDs/PropertyList-1.0.dtd">
<plist version="1.0">
<dict>
    <key>DisableSocialLogin</key>
    <true/>
    <key>DisableAutoUpdates</key>
    <string>major</string>
</dict>
</plist>
```

Windows

Set the registry values at HKLM\SOFTWARE\Policies\Amazon\Quick. The Policies hive is not user-editable and survives application updates.

The following example shows both policies configured in the registry:

```
HKLM\SOFTWARE\Policies\Amazon\Quick
  DisableSocialLogin (REG_DWORD) = 1
  DisableAutoUpdates (REG_SZ) = major
```

For MDM-specific deployment steps, see your provider's documentation:

- [macOS preference file profiles](#) on the Microsoft Intune website
- [Application & Custom Settings](#) on the Jamf Pro website

Proxy support

Amazon Quick respects the operating system's proxy configuration (system proxy on Windows, Proxy Auto-Configuration on macOS). All application network traffic, including child processes, routes through the configured enterprise proxy automatically. You don't need to configure the proxy separately.

Application data paths

The following table lists the application data paths for each operating system.

OS	Path
macOS	~/ .quickwork/
Windows	%USERPROFILE%\ .quickwork\

This directory contains local user-specific settings and configurations.

Silent installation and fleet deployment

To distribute the application silently to a managed fleet with a mobile device management (MDM) solution – including per-machine packaging for Microsoft Intune and other MDM tools, certificate

authority trust, and troubleshooting – see [Deploying Amazon Quick on desktop to a managed fleet with MDM](#).

For a single-user silent installation on Windows, the installer also supports silent mode. This installs Amazon Quick for the current user in %LOCALAPPDATA% and creates desktop and Start menu shortcuts.

```
.\Amazon-Quick-Setup.exe /S
```

Application updates

By default, Amazon Quick delivers updates automatically over HTTPS. Updates are code-signed and apply for each user on the next restart. On macOS, updates use Apple notarization. On Windows, updates use Authenticode signing.

To take control of update delivery, set the `DisableAutoUpdates` policy described in [Managed policies](#). When you block updates, you become responsible for delivering new versions to the fleet, using the same tools you use to deploy other managed applications.

Even when you block updates, Amazon Quick can still notify your users of a required update and restrict functionality remotely if a critical issue is found.

Important

We strongly recommend that you continue to allow patch updates. Patch updates deliver security fixes, performance improvements, and bug fixes. To control larger updates while your fleet still receives patches, set `DisableAutoUpdates` to `minor` or `major` instead of `patch`.

Topics

- [Set up enterprise sign-in with Microsoft Entra ID for Amazon Quick on desktop](#)
- [Set up enterprise sign-in with Google Workspace for Amazon Quick on desktop](#)
- [Set up enterprise sign-in with Okta for Amazon Quick on desktop](#)
- [Set up enterprise sign-in with Ping Identity for Amazon Quick on desktop](#)
- [Troubleshooting enterprise sign-in for Amazon Quick on desktop](#)

- [Deploying Amazon Quick on desktop to a managed fleet with MDM](#)

Set up enterprise sign-in with Microsoft Entra ID for Amazon Quick on desktop

Applies to: Enterprise Edition

Intended audience: System administrators

This page walks you through all four steps to set up enterprise sign-in with Microsoft Entra ID:

1. Create an OIDC application in Microsoft Entra ID and record its values.
2. Add the extension access in the Amazon Quick administration console.
3. Create the extension in the Amazon Quick console.
4. Download, verify, and distribute the desktop application.

Note

Before you begin, review the prerequisites in [Setting up Amazon Quick on desktop for enterprise deployments](#).

Step 1: Create an OIDC application in Microsoft Entra ID

Register a public OIDC client application in Microsoft Entra ID. The Amazon Quick desktop application uses this client to authenticate users through the authorization code flow with PKCE. This client requires no client secret. Grant the `offline_access` scope so that the application can issue refresh tokens. Without it, users must re-authenticate frequently.

For more information, see [Register an application](#) in the Microsoft Entra documentation.

To create the Entra ID app registration

1. In the Azure portal, navigate to **Microsoft Entra ID** → **App registrations** → **New registration**.

2. Configure the following settings:

Setting	Value
Name	Amazon Quick Desktop
Supported account types	Accounts in this organizational directory only (Single tenant)
Redirect URI platform	Public client/native (mobile & desktop)
Redirect URI	http://localhost:18080

3. Choose **Register**.
4. On the **Overview** page, note the **Application (client) ID** and **Directory (tenant) ID**. You need these values in later steps.

This is a public client registration. Entra ID enforces PKCE automatically for public clients.

To configure API permissions

1. In the app registration, navigate to **API permissions** → **Add a permission** → **Microsoft Graph** → **Delegated permissions**.
2. Add the following permissions: openid, email, profile, offline_access.
3. Choose **Add permissions**.
4. If your organization requires it, choose **Grant admin consent for [your organization]**.

To configure authentication settings

1. In the app registration, navigate to **Authentication**.
2. Under **Advanced settings**, leave **Allow public client flows** set to **No**.
3. Verify that http://localhost:18080 is listed under **Mobile and desktop applications**.
4. Choose **Save**.

To configure token claims

1. In the app registration, navigate to **Token configuration**.

- 2. Choose **Add optional claim**.
- 3. Select token type: **ID**.
- 4. Select the email claim and choose **Add**.

 Important

This step is required. Without the email optional claim, Microsoft Entra ID does not include the user's email address in the ID token, and Amazon Quick cannot map the token to a user. Additionally, each user who signs in must have their **Mail** attribute populated in their Entra ID profile (under **Contact Information**). The User Principal Name (UPN) alone is not sufficient — the Mail attribute must contain a value.

Record the following OIDC endpoints. You enter these values, together with the **Application (client) ID** you noted earlier, in Step 2. Replace <TENANT_ID> with your Directory (tenant) ID.

 Important

The Issuer URL must include the /v2.0 path suffix. Do not use the **Authority URL** shown in the Entra ID Endpoints panel, which omits this suffix. If the /v2.0 suffix is missing, token validation fails with an "Invalid issuer" error at sign-in time.

Field	Value
Client ID	The Application (client) ID from the app registration Overview page
Issuer URL	https://login.microsoftonline.com/<TENANT_ID>/v2.0
Authorization endpoint	https://login.microsoftonline.com/<TENANT_ID>/oauth2/v2.0/authorize
Token endpoint	https://login.microsoftonline.com/<TENANT_ID>/oauth2/v2.0/token

Field	Value
JWKS URI	<code>https://login.microsoftonline.com/<TENANT_ID>/discovery/v2.0/keys</code>

 Tip

The JWKS URI is not displayed in the Microsoft Entra ID **Endpoints** panel. You can find it by opening the **OpenID Connect metadata document** URL from the Endpoints panel and locating the `jwks_uri` field in the JSON response. Alternatively, construct it using the format shown in the preceding table.

 Important

If your application has custom signing keys because of the claims-mapping feature, you must append an `appid` query parameter containing the **Application (client) ID** to the JWKS URI. Otherwise, token validation fails because the default keys endpoint does not return the application's custom signing key. For example:

`https://login.microsoftonline.com/<TENANT_ID>/discovery/v2.0/keys?appid=<CLIENT_ID>`

To confirm the correct value, open the OpenID Connect metadata document with the same `appid` parameter appended (`https://login.microsoftonline.com/<TENANT_ID>/.well-known/openid-configuration?appid=<CLIENT_ID>`) and use the `jwks_uri` it returns. For more information, see [Validate the signature](#) in the Microsoft Entra documentation.

Step 2: Add the extension access in the Amazon Quick administration console

In the Amazon Quick administration console, add an extension access using the OIDC endpoint values and Client ID that you recorded in Step 1.

To add the extension access

1. Sign in to the Amazon Quick administration console and choose **Manage account**.
2. In the left navigation pane, under **Permissions**, choose **Extension access**.

3. Choose **Add extension access**.
4. Under **Select Service**, select **Amazon Quick (Desktop application for Quick)**, and then choose **Next**.
5. Enter the extension details using the values you recorded in Step 1:

Field	Value
Name	A name for this extension access (for example, QuickDesktop-access). This is an internal reference only and is not configured in your IdP. Use alphanumeric characters and hyphens only, with no spaces.
Description	(Optional) A description of this extension access, for your reference only.
Issuer URL	The exact OIDC issuer URL from Step 1, including any required path suffix.
Authorization Endpoint	The OIDC authorization endpoint URL from Step 1
Token Endpoint	The OIDC token endpoint URL from Step 1
JWKS URI	The JSON Web Key Set URI from Step 1
Client ID	The OIDC client identifier (Client ID or Application (client) ID) that you recorded in Step 1. Required.

6. Choose **Add**.

 Important

Verify that all values are correct before you choose **Add**. The extension access configuration cannot be edited after creation. If any value is incorrect, you must delete the extension access and create a new one.

Step 3: Create the extension in the Amazon Quick console

After you add the extension access in the Amazon Quick administration console, create the extension on the **Extensions** page in the Amazon Quick console.

To create the extension

1. In the Amazon Quick console, in the left navigation pane, choose **Extensions**. If you don't see **Extensions**, choose **More** to find it.
2. Choose **Add extension**.
3. Select the **Desktop application for Quick** extension access that you created in Step 2, and then choose **Next**.
4. Choose **Create**.

Important

Both Step 2 and Step 3 are required. If you add the extension access but do not create the extension, enterprise sign-in is not available and users see the error: "Enterprise sign-in for Quick Desktop has not been configured for this account."

Note

Creating the extension is a one-time, account-level action. After an administrator creates the extension, enterprise sign-in is available for all users in the account. Individual users do not need to enable the extension themselves — they only need to download the desktop application and sign in.

Step 4: Download, verify, and distribute the desktop application

In this step, you download and install the desktop application, verify that enterprise sign-in works, and then distribute the application to your users.

First, download the application from the Amazon Quick console.

To download the desktop application

1. In the Amazon Quick console, in the left navigation pane, choose **Extensions**. If you don't see **Extensions**, choose **More** to find it.
2. Select the Quick Desktop extension that you created in Step 3.
3. Choose the more options icon (...) for the extension.
4. Choose **Download for Windows** or **Download for Mac**, depending on your operating system.

Then, install the application. For installation instructions, see [Getting started](#).

After you install the application, verify that enterprise sign-in works.

To verify enterprise sign-in

1. Open the Amazon Quick desktop application.
2. On the sign-in screen, choose **Continue with SSO**.
3. (Optional) Select your AWS Region from the list, or choose **Dynamic** to have the application detect your Region automatically.
4. Authenticate with your corporate credentials. The application redirects to your identity provider, and then returns to the Home screen after authentication succeeds.

Tip

If sign-in fails, verify the values you entered in Step 2 against the OIDC endpoints and Client ID from Step 1. If any value is incorrect, delete the extension access under **Permissions** → **Extension access** in the Amazon Quick administration console, and repeat Step 2 with the correct values. For more help, see [Troubleshooting enterprise sign-in for Amazon Quick on desktop](#).

Finally, after you verify the setup, distribute the application to your users. Direct them to [Getting started](#) for download, installation, and sign-in instructions. Users choose **Continue with SSO** and sign in with their corporate credentials — no additional per-user configuration is required.

Set up enterprise sign-in with Google Workspace for Amazon Quick on desktop

Applies to: Enterprise Edition

Intended audience: System administrators

This page walks you through all four steps to set up enterprise sign-in with Google Workspace:

1. Create an OIDC application in Google Workspace and record its values.
2. Add the extension access in the Amazon Quick administration console.
3. Create the extension in the Amazon Quick console.
4. Download, verify, and distribute the desktop application.

Note

Before you begin, review the prerequisites in [Setting up Amazon Quick on desktop for enterprise deployments](#).

Step 1: Create an OIDC application in Google Workspace

Register an OAuth client in the Google Cloud Console. The Amazon Quick desktop application uses this client to authenticate users through the authorization code flow with PKCE. Google issues a refresh token automatically when the authorization request includes the `access_type=offline` parameter. Quick includes this parameter automatically. No additional scope or grant type configuration is required.

For more information, see [OAuth 2.0 for Mobile & Desktop Apps](#) in the Google for Developers documentation.

To create the Google OAuth client

1. In the Google Cloud Console, navigate to **APIs & Services** → **Credentials** → **Create Credentials** → **OAuth client ID**.

2. Configure the following settings:

Setting	Value
Application type	Desktop app
Name	Amazon Quick Desktop

3. Choose **Create**.

4. Note the **Client ID** and **Client secret**. You need these values in later steps.

To configure the OAuth consent screen

1. In the Google Cloud Console, navigate to **Google Auth Platform** → **Branding**.
2. Set the **User type** to **Internal**. This restricts sign-in to users in your Google Workspace organization.
3. Fill in the required app information and choose **Save**.

To configure scopes

1. In the Google Cloud Console, navigate to **Google Auth Platform** → **Data Access**.
2. Add the following scopes: openid, email, profile.
3. Choose **Save**.

Record the following OIDC endpoints. You enter these values, together with the **Client ID** you noted earlier, in Step 2.

Field	Value
Client ID	The Client ID from your OAuth client
Issuer URL	<code>https://accounts.google.com</code>
Authorization endpoint	<code>https://accounts.google.com/o/oauth2/v2/auth</code>

Field	Value
Token endpoint	<code>https://oauth2.googleapis.com/token?client_secret=<CLIENT_SECRET></code>
JWKS URI	<code>https://www.googleapis.com/oauth2/v3/certs</code>

 Note

Append your client secret to the token endpoint as a `client_secret` query parameter so that token exchange succeeds—for example, `https://oauth2.googleapis.com/token?client_secret=<CLIENT_SECRET>`. Replace `<CLIENT_SECRET>` with the client secret generated for your OAuth client.

Step 2: Add the extension access in the Amazon Quick administration console

In the Amazon Quick administration console, add an extension access using the OIDC endpoint values and Client ID that you recorded in Step 1.

To add the extension access

1. Sign in to the Amazon Quick administration console and choose **Manage account**.
2. In the left navigation pane, under **Permissions**, choose **Extension access**.
3. Choose **Add extension access**.
4. Under **Select Service**, select **Amazon Quick (Desktop application for Quick)**, and then choose **Next**.
5. Enter the extension details using the values you recorded in Step 1:

Field	Value
Name	A name for this extension access (for example, QuickDesktop-access). This is an internal reference only and is not configured in your IdP. Use alphanumeric characters and hyphens only, with no spaces.

Field	Value
Description	(Optional) A description of this extension access, for your reference only.
Issuer URL	The exact OIDC issuer URL from Step 1, including any required path suffix.
Authorization Endpoint	The OIDC authorization endpoint URL from Step 1
Token Endpoint	The OIDC token endpoint URL from Step 1
JWKS URI	The JSON Web Key Set URI from Step 1
Client ID	The OIDC client identifier (Client ID or Application (client) ID) that you recorded in Step 1. Required.

6. Choose **Add**.

 Important

Verify that all values are correct before you choose **Add**. The extension access configuration cannot be edited after creation. If any value is incorrect, you must delete the extension access and create a new one.

Step 3: Create the extension in the Amazon Quick console

After you add the extension access in the Amazon Quick administration console, create the extension on the **Extensions** page in the Amazon Quick console.

To create the extension

1. In the Amazon Quick console, in the left navigation pane, choose **Extensions**. If you don't see **Extensions**, choose **More** to find it.
2. Choose **Add extension**.
3. Select the **Desktop application for Quick** extension access that you created in Step 2, and then choose **Next**.
4. Choose **Create**.

 Important

Both Step 2 and Step 3 are required. If you add the extension access but do not create the extension, enterprise sign-in is not available and users see the error: "Enterprise sign-in for Quick Desktop has not been configured for this account."

 Note

Creating the extension is a one-time, account-level action. After an administrator creates the extension, enterprise sign-in is available for all users in the account. Individual users do not need to enable the extension themselves — they only need to download the desktop application and sign in.

Step 4: Download, verify, and distribute the desktop application

In this step, you download and install the desktop application, verify that enterprise sign-in works, and then distribute the application to your users.

First, download the application from the Amazon Quick console.

To download the desktop application

1. In the Amazon Quick console, in the left navigation pane, choose **Extensions**. If you don't see **Extensions**, choose **More** to find it.
2. Select the Quick Desktop extension that you created in Step 3.
3. Choose the more options icon (...) for the extension.
4. Choose **Download for Windows** or **Download for Mac**, depending on your operating system.

Then, install the application. For installation instructions, see [Getting started](#).

After you install the application, verify that enterprise sign-in works.

To verify enterprise sign-in

1. Open the Amazon Quick desktop application.
2. On the sign-in screen, choose **Continue with SSO**.

3. (Optional) Select your AWS Region from the list, or choose **Dynamic** to have the application detect your Region automatically.
4. Authenticate with your corporate credentials. The application redirects to your identity provider, and then returns to the Home screen after authentication succeeds.

Tip

If sign-in fails, verify the values you entered in Step 2 against the OIDC endpoints and Client ID from Step 1. If any value is incorrect, delete the extension access under **Permissions** → **Extension access** in the Amazon Quick administration console, and repeat Step 2 with the correct values. For more help, see [Troubleshooting enterprise sign-in for Amazon Quick on desktop](#).

Finally, after you verify the setup, distribute the application to your users. Direct them to [Getting started](#) for download, installation, and sign-in instructions. Users choose **Continue with SSO** and sign in with their corporate credentials — no additional per-user configuration is required.

Set up enterprise sign-in with Okta for Amazon Quick on desktop

Applies to: Enterprise Edition

Intended audience: System administrators

This page walks you through all four steps to set up enterprise sign-in with Okta:

1. Create an OIDC application in Okta and record its values.
2. Add the extension access in the Amazon Quick administration console.
3. Create the extension in the Amazon Quick console.
4. Download, verify, and distribute the desktop application.

Note

Before you begin, review the prerequisites in [Setting up Amazon Quick on desktop for enterprise deployments](#).

Step 1: Create an OIDC application in Okta

Register an OIDC native application in Okta. The Amazon Quick desktop application uses this client to authenticate users through the authorization code flow with PKCE, and requests the `openid`, `profile`, `email`, and `offline_access` scopes at sign-in.

For more information, see [Create OpenID Connect app integrations](#) in the Okta documentation.

To create the Okta OIDC Native Application

1. In the Okta Admin Console, navigate to **Applications** → **Applications** → **Create App Integration**.
2. Select **OIDC - OpenID Connect** as the sign-in method.
3. Select **Native Application** as the application type, then choose **Next**.
4. Configure the following settings:

Setting	Value
App integration name	Amazon Quick Desktop
Grant type	Authorization Code and Refresh Token
Sign-in redirect URIs	<code>http://localhost:18080</code>
Assignments	Assign to the appropriate users or groups

5. Choose **Save**.
6. On the **General** tab, note the **Client ID**.

Okta enforces PKCE (S256) automatically for native applications.

To verify authentication settings

1. In the app integration, go to the **General** tab.
2. Under **General Settings**, confirm that the application type is **Native**, client authentication is **None** (public client), and PKCE is **Required**.
3. Under **LOGIN**, confirm that `http://localhost:18080` is listed as a redirect URI.
4. Choose **Save** if you made any changes.

Record the following OIDC endpoints. You enter these values, together with the **Client ID** you noted earlier, in Step 2. Replace `<OKTA_DOMAIN>` with your Okta domain (for example, `your-org.okta.com`). For the endpoint reference, see [Authorization Code with PKCE](#) in the Okta documentation.

Field	Value
Client ID	The Client ID from the General tab
Issuer URL	<code>https://<OKTA_DOMAIN></code>
Authorization endpoint	<code>https://<OKTA_DOMAIN>/oauth2/v1/authorize</code>
Token endpoint	<code>https://<OKTA_DOMAIN>/oauth2/v1/token</code>
JWKS URI	<code>https://<OKTA_DOMAIN>/oauth2/v1/keys</code>

Use the standard Okta domain, not the admin subdomain

Use your standard Okta domain (for example, `your-org.okta.com`) in these endpoints, not the `-admin` subdomain (`your-org-admin.okta.com`) shown in the address bar of the admin console. The `-admin` subdomain does not host the OAuth endpoints, so using it causes sign-in to fail with an HTTP 404 error.

Step 2: Add the extension access in the Amazon Quick administration console

In the Amazon Quick administration console, add an extension access using the OIDC endpoint values and Client ID that you recorded in Step 1.

To add the extension access

1. Sign in to the Amazon Quick administration console and choose **Manage account**.
2. In the left navigation pane, under **Permissions**, choose **Extension access**.
3. Choose **Add extension access**.
4. Under **Select Service**, select **Amazon Quick (Desktop application for Quick)**, and then choose **Next**.
5. Enter the extension details using the values you recorded in Step 1:

Field	Value
Name	A name for this extension access (for example, QuickDesktop-access). This is an internal reference only and is not configured in your IdP. Use alphanumeric characters and hyphens only, with no spaces.
Description	(Optional) A description of this extension access, for your reference only.
Issuer URL	The exact OIDC issuer URL from Step 1, including any required path suffix.
Authorization Endpoint	The OIDC authorization endpoint URL from Step 1
Token Endpoint	The OIDC token endpoint URL from Step 1
JWKS URI	The JSON Web Key Set URI from Step 1
Client ID	The OIDC client identifier (Client ID or Application (client) ID) that you recorded in Step 1. Required.

6. Choose **Add**.

Important

Verify that all values are correct before you choose **Add**. The extension access configuration cannot be edited after creation. If any value is incorrect, you must delete the extension access and create a new one.

Step 3: Create the extension in the Amazon Quick console

After you add the extension access in the Amazon Quick administration console, create the extension on the **Extensions** page in the Amazon Quick console.

To create the extension

1. In the Amazon Quick console, in the left navigation pane, choose **Extensions**. If you don't see **Extensions**, choose **More** to find it.
2. Choose **Add extension**.
3. Select the **Desktop application for Quick** extension access that you created in Step 2, and then choose **Next**.
4. Choose **Create**.

Important

Both Step 2 and Step 3 are required. If you add the extension access but do not create the extension, enterprise sign-in is not available and users see the error: "Enterprise sign-in for Quick Desktop has not been configured for this account."

Note

Creating the extension is a one-time, account-level action. After an administrator creates the extension, enterprise sign-in is available for all users in the account. Individual users do not need to enable the extension themselves — they only need to download the desktop application and sign in.

Step 4: Download, verify, and distribute the desktop application

In this step, you download and install the desktop application, verify that enterprise sign-in works, and then distribute the application to your users.

First, download the application from the Amazon Quick console.

To download the desktop application

1. In the Amazon Quick console, in the left navigation pane, choose **Extensions**. If you don't see **Extensions**, choose **More** to find it.
2. Select the Quick Desktop extension that you created in Step 3.
3. Choose the more options icon (...) for the extension.
4. Choose **Download for Windows** or **Download for Mac**, depending on your operating system.

Then, install the application. For installation instructions, see [Getting started](#).

After you install the application, verify that enterprise sign-in works.

To verify enterprise sign-in

1. Open the Amazon Quick desktop application.
2. On the sign-in screen, choose **Continue with SSO**.
3. (Optional) Select your AWS Region from the list, or choose **Dynamic** to have the application detect your Region automatically.
4. Authenticate with your corporate credentials. The application redirects to your identity provider, and then returns to the Home screen after authentication succeeds.

Tip

If sign-in fails, verify the values you entered in Step 2 against the OIDC endpoints and Client ID from Step 1. If any value is incorrect, delete the extension access under **Permissions** → **Extension access** in the Amazon Quick administration console, and repeat Step 2 with the correct values. For more help, see [Troubleshooting enterprise sign-in for Amazon Quick on desktop](#).

Finally, after you verify the setup, distribute the application to your users. Direct them to [Getting started](#) for download, installation, and sign-in instructions. Users choose **Continue with SSO** and sign in with their corporate credentials — no additional per-user configuration is required.

Set up enterprise sign-in with Ping Identity for Amazon Quick on desktop

Applies to: Enterprise Edition

Intended audience: System administrators

This page walks you through all four steps to set up enterprise sign-in with Ping Identity:

1. Create an OIDC application in your Ping Identity product (PingFederate or PingOne) and record its values.
2. Add the extension access in the Amazon Quick administration console.
3. Create the extension in the Amazon Quick console.
4. Download, verify, and distribute the desktop application.

Note

Before you begin, review the prerequisites in [Setting up Amazon Quick on desktop for enterprise deployments](#).

Step 1: Create an OIDC application in Ping Identity

Register a public OIDC client in your Ping Identity product. The Amazon Quick desktop application uses this client to authenticate users through the authorization code flow with PKCE. This client requires no client secret. Enable the Refresh Token grant type and grant the `offline_access` scope. Choose the instructions for your Ping Identity product.

PingFederate

For more information, see [Setting up an OIDC application in PingFederate](#) in the Ping Identity documentation.

To create the PingFederate OIDC client

1. In the PingFederate administrative console, go to **Applications** → **OAuth** → **Clients**, and choose **Add Client**.
2. In the **Client ID** field, enter a unique identifier for this client.
3. In the **Name** field, enter Amazon Quick Desktop.
4. For **Client Authentication**, select **None**.
5. In the **Redirection URI** section, enter `http://localhost:18080` and choose **Add**.
6. In the **Allowed Grant Types** list, select **Authorization Code** and **Refresh Token**.
7. Select the **Require Proof Key for Code Exchange (PKCE)** checkbox.
8. Under **Common Scopes**, grant the following: `openid`, `email`, `profile`, `offline_access`.
9. Choose **Save**.
10. Note the **Client ID**. You need this value in later steps.

To configure the OIDC policy

1. In the PingFederate administrative console, go to **Applications** → **OAuth** → **OpenID Connect Policy Management**.
2. Select the OIDC policy associated with this client, or choose **Add Policy** to create one.
3. Select the **Return ID Token On Refresh Grant** checkbox. This ensures that the desktop application receives a fresh ID token with current claims when refreshing the session.
4. Under **Attribute Contract**, verify that the `email` claim is included and mapped to the corresponding user attribute in your authentication source. The `email` claim must be present in tokens issued during both initial authentication and refresh token grants.
5. Choose **Save**.

Record the following OIDC endpoints. You enter these values, together with the **Client ID** you noted earlier, in Step 2. Replace `<PINGFEDERATE_HOST>` with your PingFederate server hostname.

Field	Value
Client ID	The Client ID you entered when you created the client
Issuer URL	<code>https://<PINGFEDERATE_HOST></code>

Field	Value
Authorization endpoint	https://<PINGFEDERATE_HOST>/as/authorization.oauth2
Token endpoint	https://<PINGFEDERATE_HOST>/as/token.oauth2
JWKS URI	https://<PINGFEDERATE_HOST>/pf/JWKS

PingOne

For more information, see [Editing an application – Native](#) in the Ping Identity documentation.

To create the PingOne OIDC native application

1. In the PingOne admin console, go to **Applications** → **Applications** and choose the **+** icon.
2. Enter Amazon Quick Desktop as the application name.
3. In the **Application Type** section, select **Native**, then choose **Save**.
4. On the **Configuration** tab, choose **Edit** and configure the following settings:

Setting	Value
Response Type	Code
Grant Type	Authorization Code and Refresh Token
PKCE Enforcement	S256
Redirect URIs	http://localhost:18080
Token Endpoint Authentication Method	None

5. Choose **Save**.
6. On the **Resources** tab, add the following scopes: openid, email, profile, offline_access.

⚠ Add scopes on the Resources tab, not the Configuration tab

This is a separate step from the **Configuration** tab. PingOne grants only the `openid` scope by default, so you must explicitly add `email`, `profile`, and `offline_access` on the **Resources** tab. If you skip this step, sign-in fails with "User info request failed (HTTP 401)". This error occurs because sign-in requires tokens issued with these scopes, and PingOne does not grant them unless you add them on this tab.

7. On the **Attribute Mappings** tab, verify that the `email` attribute is mapped to the user's email address.
8. Toggle the application to **Enabled**.
9. Note the **Client ID** and **Environment ID** from the **Configuration** tab.

Note

The PingOne domain varies by region. The following examples use `.com`. Replace the domain with the one for your environment (for example, `.ca`, `.eu`, or `.asia`).

Record the following OIDC endpoints. You enter these values, together with the **Client ID** you noted earlier, in Step 2. Replace `<ENV_ID>` with your PingOne environment ID.

Field	Value
Client ID	The Client ID from the Configuration tab
Issuer URL	<code>https://auth.pingone.com/<ENV_ID>/as</code>
Authorization endpoint	<code>https://auth.pingone.com/<ENV_ID>/as/authorize</code>
Token endpoint	<code>https://auth.pingone.com/<ENV_ID>/as/token</code>
JWKS URI	<code>https://auth.pingone.com/<ENV_ID>/as/jwks</code>

⚠ Use the correct JWKS URI path for PingOne

For the JSON Web Key Set (JWKS) URI, use the `/as/jwks` path shown in the preceding table. Do not use `/.well-known/jwks.json`, which can appear as a placeholder in some PingOne forms. PingOne does not use this path, and sign-in token validation fails if you enter it.

Step 2: Add the extension access in the Amazon Quick administration console

In the Amazon Quick administration console, add an extension access using the OIDC endpoint values and Client ID that you recorded in Step 1.

To add the extension access

1. Sign in to the Amazon Quick administration console and choose **Manage account**.
2. In the left navigation pane, under **Permissions**, choose **Extension access**.
3. Choose **Add extension access**.
4. Under **Select Service**, select **Amazon Quick (Desktop application for Quick)**, and then choose **Next**.
5. Enter the extension details using the values you recorded in Step 1:

Field	Value
Name	A name for this extension access (for example, QuickDesktop-access). This is an internal reference only and is not configured in your IdP. Use alphanumeric characters and hyphens only, with no spaces.
Description	(Optional) A description of this extension access, for your reference only.
Issuer URL	The exact OIDC issuer URL from Step 1, including any required path suffix.
Authorization Endpoint	The OIDC authorization endpoint URL from Step 1

Field	Value
Token Endpoint	The OIDC token endpoint URL from Step 1
JWKS URI	The JSON Web Key Set URI from Step 1
Client ID	The OIDC client identifier (Client ID or Application (client) ID) that you recorded in Step 1. Required.

6. Choose **Add**.

 Important

Verify that all values are correct before you choose **Add**. The extension access configuration cannot be edited after creation. If any value is incorrect, you must delete the extension access and create a new one.

Step 3: Create the extension in the Amazon Quick console

After you add the extension access in the Amazon Quick administration console, create the extension on the **Extensions** page in the Amazon Quick console.

To create the extension

1. In the Amazon Quick console, in the left navigation pane, choose **Extensions**. If you don't see **Extensions**, choose **More** to find it.
2. Choose **Add extension**.
3. Select the **Desktop application for Quick** extension access that you created in Step 2, and then choose **Next**.
4. Choose **Create**.

 Important

Both Step 2 and Step 3 are required. If you add the extension access but do not create the extension, enterprise sign-in is not available and users see the error: "Enterprise sign-in for Quick Desktop has not been configured for this account."

Note

Creating the extension is a one-time, account-level action. After an administrator creates the extension, enterprise sign-in is available for all users in the account. Individual users do not need to enable the extension themselves — they only need to download the desktop application and sign in.

Step 4: Download, verify, and distribute the desktop application

In this step, you download and install the desktop application, verify that enterprise sign-in works, and then distribute the application to your users.

First, download the application from the Amazon Quick console.

To download the desktop application

1. In the Amazon Quick console, in the left navigation pane, choose **Extensions**. If you don't see **Extensions**, choose **More** to find it.
2. Select the Quick Desktop extension that you created in Step 3.
3. Choose the more options icon (...) for the extension.
4. Choose **Download for Windows** or **Download for Mac**, depending on your operating system.

Then, install the application. For installation instructions, see [Getting started](#).

After you install the application, verify that enterprise sign-in works.

To verify enterprise sign-in

1. Open the Amazon Quick desktop application.
2. On the sign-in screen, choose **Continue with SSO**.
3. (Optional) Select your AWS Region from the list, or choose **Dynamic** to have the application detect your Region automatically.
4. Authenticate with your corporate credentials. The application redirects to your identity provider, and then returns to the Home screen after authentication succeeds.

Tip

If sign-in fails, verify the values you entered in Step 2 against the OIDC endpoints and Client ID from Step 1. If any value is incorrect, delete the extension access under **Permissions** → **Extension access** in the Amazon Quick administration console, and repeat Step 2 with the correct values. For more help, see [Troubleshooting enterprise sign-in for Amazon Quick on desktop](#).

Finally, after you verify the setup, distribute the application to your users. Direct them to [Getting started](#) for download, installation, and sign-in instructions. Users choose **Continue with SSO** and sign in with their corporate credentials — no additional per-user configuration is required.

Troubleshooting enterprise sign-in for Amazon Quick on desktop

Applies to: Enterprise Edition

Intended audience: System administrators

Use the following guidance to resolve common enterprise sign-in issues, regardless of which identity provider you use.

Tip

To help diagnose a sign-in problem, you can export application logs from the sign-in screen. Include these logs when you contact your administrator or AWS Support.

Note

If the application cannot reach the sign-in page, complete authentication, or load content, the issue might be network related. In restricted environments, confirm that the required domains are on your allow list and that firewall and VPN settings are not blocking the connections. For the list of required domains, see [Network access and required domains](#).

redirect_mismatch error

Verify that the redirect URI in your IdP is exactly `http://localhost:18080` and is configured as a public client or native platform.

Access blocked: app not verified (Google Workspace)

This error means the OAuth consent screen is set to **External**. In the Google Cloud Console, navigate to **Google Auth Platform** → **Branding** and set the **User type** to **Internal**, which restricts sign-in to users in your Google Workspace organization.

User not found after sign-in

This error has two common causes:

1. **The email claim is not being returned in the token.** For Microsoft Entra ID, you must add the `email` optional claim to the ID token under Token configuration (see Step 1). Additionally, the user's **Mail** attribute must be populated in their Entra ID profile. The User Principal Name (UPN) alone is not sufficient.
2. **No matching user exists in Amazon Quick.** The email in the token must exactly match the email of a provisioned user. For IAM Identity Center accounts, verify the user's email in Identity Center matches. Email matching is case-sensitive.

Token validation failure

This error has two common causes:

1. **The issuer URL does not match.** Verify that the issuer URL in the extension access configuration matches the issuer URL in your IdP's OIDC configuration exactly.
2. **The ID token is signed with an unsupported algorithm.** Amazon Quick accepts asymmetric signatures only. If your IdP signs ID tokens with HMAC (HS256, HS384, or HS512), validation fails regardless of the JWKS URI you configured, because the signing key is never published in a JWKS. Reconfigure your IdP to use a supported algorithm. For the list, see [Supported token signing algorithms](#).

Invalid issuer error (Microsoft Entra ID)

If sign-in fails with "Invalid issuer: `https://login.microsoftonline.com/TENANT_ID/v2.0`", verify that the Issuer URL in your extension access configuration includes the `/v2.0` path suffix. The Entra ID v2.0 endpoint issues tokens with an `iss` claim that includes `/v2.0`. If the suffix is missing, delete the extension access and recreate it with the correct Issuer URL.

Enterprise sign-in not configured for this account

This error means the extension access was created but the extension itself was not. In the Amazon Quick console, in the left navigation pane, choose **Extensions** (you might need to choose **More** to find it), and create the extension, selecting the extension access you previously configured.

User info request failed (HTTP 504)

This is a transient backend timeout. Sign in to your Amazon Quick account via the web browser first, then retry the desktop sign-in. If the error persists, verify network connectivity to the Amazon Quick service endpoint. For the list of required domains, see [Network access and required domains](#).

Consent or permission errors (Microsoft Entra ID)

Grant admin consent for the required API permissions in the Azure portal. Navigate to the app registration's **API permissions** page and choose **Grant admin consent for [your organization]**.

Session expires frequently

Verify that your IdP is configured to issue refresh tokens. For Microsoft Entra ID, the `offline_access` scope is required. For Google Workspace, include `access_type=offline` in the authorization request (handled automatically by Quick). For Okta, the Refresh Token grant type must be enabled and the `offline_access` scope must be granted. For Ping Identity, the Refresh Token grant type must be enabled and the `offline_access` scope must be granted. For PingFederate, also verify that **Return ID Token On Refresh Grant** is selected in the OIDC policy.

Application not enabled (PingOne)

If authentication fails immediately without reaching the PingOne login page, verify that the application toggle is set to **Enabled** in the PingOne admin console.

User info request failed (HTTP 401) (PingOne)

The required scopes are not enabled on the application's **Resources** tab. This is a separate step from the **Configuration** tab: PingOne grants only the `openid` scope by default. In the PingOne admin console, open the application's **Resources** tab and add the `email`, `profile`, and `offline_access` scopes.

Token validation failure with a PingOne JSON Web Key Set (JWKS) URI

Verify that the JWKS URI in your extension access configuration uses the `/as/jwks` path (for example, `https://auth.pingone.com/<ENV_ID>/as/jwks`). Do not use `/.well-known/jwks.json`, which can appear as a placeholder in some PingOne forms. PingOne does not use this path.

Missing email claim after refresh (PingFederate)

Verify that the `email` claim is included in the OIDC policy **Attribute Contract** and mapped to the correct user attribute. The mapping must produce the `email` claim for both initial authentication and refresh token grants.

Deploying Amazon Quick on desktop to a managed fleet with MDM

Applies to: Enterprise Edition

Intended audience: System administrators

As an administrator, you can distribute the Amazon Quick desktop application to a managed fleet with mobile device management (MDM), without any per-machine interaction. This page describes the deployment using Microsoft Intune for both macOS and Windows. The artifacts it describes – a signed macOS `.pkg` and a per-machine Windows `.exe` – also work with other MDM solutions, such as Jamf, Kandji, and VMware Workspace ONE. Only the upload mechanics differ between solutions.

This page describes how to complete the following tasks:

1. Get the per-machine installer artifacts from the download service.
2. Deploy the macOS application with Microsoft Intune.
3. Deploy the Windows application with Microsoft Intune.
4. (Optional) Trust a corporate certificate authority for networks that inspect encrypted traffic.

Note

This page covers distributing the application binary. To let users sign in with their corporate credentials, you must also configure enterprise single sign-on. For more information, see [Setting up Amazon Quick on desktop for enterprise deployments](#). To enforce enterprise SSO as the only sign-in method by disabling consumer social login across your fleet, see [Managed policies](#).

Prerequisites

Before you begin, verify that you have the following:

- Administrator access to your MDM solution, and a fleet of devices that are enrolled in it.
- Devices that meet the Amazon Quick desktop application system requirements. For the macOS and Windows requirements, see [Prerequisites](#).
- For Windows, devices that run Windows Pro, Enterprise, or Education. The application deployment method for Windows uses the Win32 app type, which Microsoft Intune does not support on Windows Home. MDM enrollment also requires a Windows client edition (Windows Server is not supported for MDM enrollment).
- For Windows, the Microsoft Win32 Content Prep Tool (IntuneWinAppUtil.exe), which packages the installer into the .intunewin format that the Win32 app type requires. For more information, see [Microsoft Win32 Content Prep Tool](#) on GitHub.

⚠ Use per-machine installers only

Use the per-machine installer variants described in this page, not the per-user installers that your users download individually. An MDM solution runs as the system account with no signed-in user, so a per-user installer does not deploy correctly in that context.

Get the installer artifacts

The Amazon Quick download service at `desktop.downloads.quick.aws.com` provides the desktop installers. For MDM deployment, use the per-machine installer for each platform. The per-machine installer differs from the per-user installer that individual users download from the standard download page:

- **macOS** – a signed, notarized component .pkg, named Amazon Quick-*version*-arm64.pkg, that installs Amazon Quick.app to the /Applications folder as the root user.
- **Windows** – a per-machine .exe, named Amazon Quick-Admin Setup *version*.exe, that installs the application for all users into C:\Program Files\Amazon Quick\.

Each platform publishes a channel manifest that lists the current version and a SHA-512 checksum. Fetch the manifest to find the version to deploy and the checksum to verify your download against.

```
# macOS manifest
curl -fsSL https://desktop.downloads.quick.aws.com/darwin/arm64/quick-external-cloud-mac.yml

# Windows per-machine manifest
curl -fsSL https://desktop.downloads.quick.aws.com/win32/x64/quick-external-cloud-admin.yml
```

The manifest includes two keys that you use for MDM deployment.

version

The version of the most recent build released externally. Use this value wherever a procedure in this topic asks you for *version*.

path

The version-specific artifact, relative to the platform base URL. On Windows, you append path to `https://desktop.downloads.quick.aws.com/win32/x64/` to download the per-machine installer. On macOS, path refers to the archive that the application auto-updater consumes rather than a deployable installer, so you use version with the .pkg URL pattern that follows.

Download the per-machine artifact for each platform. In the following URL patterns, replace *version* with the version from the manifest.

```
# macOS .pkg (per-machine, installs to /Applications)
https://desktop.downloads.quick.aws.com/darwin/arm64/quick-external-cloud/Amazon
Quick-version-arm64.pkg
```



```
# Windows per-machine .exe (installs for all users into Program Files)
https://desktop.downloads.quick.aws.com/win32/x64/quick-external-cloud-admin/Amazon
Quick-Admin Setup version.exe
```

A pre-production beta channel also exists for internal testing. It updates itself automatically and does not publish a public first-time installer, so it is not used for MDM deployment. For the standard download page that individual users use, see [Downloading and installing](#).

Static and version-specific download URLs

The Windows per-machine .exe also has a static download link that resolves to the most recent build released externally, so you do not need to fetch the manifest first. A static link can trail the version in the manifest while a release finishes rolling out. If you need a specific version, or you want to verify your download against the SHA-512 checksum, fetch the manifest and use the version-specific URL instead.

```
https://desktop.downloads.quick.aws.com/windows/x64/Amazon-Quick-Admin.exe
```

The macOS .pkg has no static link, because it is published only under its version-specific channel path. Get the current version from the manifest, and then build the .pkg URL with that version.

Deploying the macOS application with Microsoft Intune

The macOS installer is a signed, notarized component .pkg that installs Amazon Quick.app to the /Applications folder.

To add and assign the macOS application

1. In the Microsoft Intune admin center, go to **Apps** → **All apps** → **Add**, and then select the **macOS app (PKG)** app type.

Note

Use the **macOS app (PKG)** type, not the line-of-business app type.

2. Choose **Select app package file**, and then upload the .pkg that you downloaded.
3. In the **Included apps** detection list, keep only the main application bundle and remove the helper bundles.

Keep only the main application bundle

The macOS installer includes the main application plus four nested helper components inside Contents/Frameworks/. When you upload the .pkg, Microsoft Intune automatically populates the **Included apps** list with all five bundles. Delete the four helper rows (the ones whose IDs end in .e0 through .e3, such as **Amazon Quick Helper (GPU)**) so that exactly one entry remains, with the following values.

Field	Value
App bundle ID	com.aws.QuickWork.mac
Build number	The numeric part of the manifest version (for example, 1.0.3377). Use the value from the manifest rather than this example.

If you leave the helper rows in place, macOS might report installation failure code 0x87D13BA2 even when the application files are present.

Note

The app bundle ID (com.aws.QuickWork.mac) is the application's CFBundleIdentifier, which Microsoft Intune uses for detection. It is intentionally different from the installer's package receipt identifier (com.amazon.QuickWork.mac), which you use with `pkgutil --forget` when you clean a test device.

4. Set **Ignore app version** to **Yes**. The application includes a built-in updater that keeps it current from the channel manifest, so let the application manage its own version. With this setting, Microsoft Intune installs the application once if the bundle ID is absent, and then does not reinstall it when the version changes. Without it, Microsoft Intune repeatedly attempts to reinstall when it detects version drift between what it deployed and what the auto-updater has since applied.
5. On the **Assignments** tab, add a device or user group with the **Required** intent.

On each assigned device, the installation runs silently and the application is installed at /Applications/Amazon Quick.app. To verify the installation on a device, run the following command and confirm that it returns com.aws.QuickWork.mac.

```
defaults read "/Applications/Amazon Quick.app/Contents/Info" CFBundleIdentifier
```

Device-side deployment logs are in the /Library/Logs/Microsoft/Intune/ directory.

Deploying the Windows application with Microsoft Intune

The Windows installer is an NSIS .exe. The Microsoft Intune Win32 app type does not accept a raw .exe, so you first wrap the installer into a .intunewin package. Use the **Windows app (Win32)** app type. The line-of-business app type accepts only .msi, .appx, and .msix files, so it does not apply.

The following table lists the key facts for the per-machine Windows installer.

Item	Value
Silent install flag	/S (a capital S, not /silent)
Silent uninstall flag	/S
Install path	C:\Program Files\Amazon Quick\
Binary name	Amazon Quick.exe
Uninstaller name	Uninstall Amazon Quick.exe
Required edition	Windows Pro, Enterprise, or Education (not Home)

To wrap the installer into a .intunewin package

1. Place the .exe installer by itself in a source folder.
2. Run the Win32 Content Prep Tool, replacing *version* with the version you downloaded.

```
IntuneWinAppUtil.exe -c C:\IntuneSource -s "Amazon Quick-Admin Setup version.exe" -  
o C:\IntuneOutput -q
```

The tool creates C:\IntuneOutput\Amazon Quick-Admin Setup *version*.intunewin.

To create and assign the Win32 application

1. In the Microsoft Intune admin center, go to **Apps** → **All apps** → **Add**, select the **Windows app (Win32)** app type, and then upload the .intunewin package.
2. On the **Program** tab, configure the following settings. Replace *version* with the version you downloaded.
 - **Install command** – "Amazon Quick-Admin Setup *version*.exe" /S
 - **Uninstall command** – "C:\Program Files\Amazon Quick\Uninstall Amazon Quick.exe" /S
 - **Install behavior** – **System**
3. On the **Requirements** tab, set the operating system architecture to **x64** and set the minimum operating system to a supported Windows client version.
4. On the **Detection rules** tab, add a manually configured rule with the following settings.

Field	Value
Rule type	File
Path	C:\Program Files\Amazon Quick
File or folder	Amazon Quick.exe
Detection method	File or folder exists
Associated with a 32-bit app on 64-bit clients	No

5. On the **Assignments** tab, add a device group with the **Required** intent.

The Intune Management Extension installs on a device the first time you assign a Win32 application, which can take a few minutes and two device check-ins. To trigger a check-in on a

device, go to **Settings** → **Accounts** → **Access work or school**, select your account, and choose **Sync**. To watch the installation on a device, tail the log with the following PowerShell command.

```
Get-Content "C:\ProgramData\Microsoft\IntuneManagementExtension\Logs\AppWorkload.log" -  
Wait -Tail 50
```

Trusting a corporate certificate authority

If your network inspects encrypted traffic or uses a custom corporate certificate authority (CA), the Amazon Quick desktop application must trust that CA. Otherwise, sign-in and responses can fail on the corporate network. Examples of traffic-inspection proxies include Zscaler and Netskope.

Deploy the CA into the operating system trust store with an MDM certificate profile. The application reads the operating system trust store when it starts, so a CA that you install there is trusted for all of the application's outbound connections. This approach works on both platforms and requires no environment variables, scripts, or certificate file placement.

macOS

Deploy a configuration profile with a **Trusted certificate** payload that targets the **System keychain**.

Windows

Deploy a **Trusted certificate** configuration profile to the computer store under **Trusted Root**.

Note

The certificate profile is the reliable path for the application itself. Deploy environment variables such as REQUESTS_CA_BUNDLE, SSL_CERT_FILE, or NODE_EXTRA_CA_CERTS only if other, non-Amazon Quick tools on the device require the CA through those variables.

Advanced: environment variables for mixed toolchains on Windows

Do not use per-user path variables in scripts that run under Intune

Do not use %HOMEDRIVE%%HOMEPATH% in scripts that run under Intune. Intune runs scripts as the system account, where those variables resolve to a service profile path (for example, C:\Windows\system32\config\systemprofile), not the user's home directory. Place certificate files and environment-variable scripts at a machine-wide path such as C:\ProgramData\Amazon Quick\certs\ instead. Use SETX /M to write machine-scoped environment variables. New processes inherit these variables after the next sign-in or restart.

The following example install script places the CA bundle at a machine-wide path and sets the environment variables.

```
@echo off
mkdir "C:\ProgramData\Amazon Quick\certs" 2>nul
copy /Y "%~dp0tls-ca-bundle.pem" "C:\ProgramData\Amazon Quick\certs\tls-ca-bundle.pem"
SETX /M REQUESTS_CA_BUNDLE "C:\ProgramData\Amazon Quick\certs\tls-ca-bundle.pem"
SETX /M SSL_CERT_FILE "C:\ProgramData\Amazon Quick\certs\tls-ca-bundle.pem"
exit /b 0
```

The NODE_EXTRA_CA_CERTS variable is not consumed by the Amazon Quick agent, so the certificate profile remains the reliable path for the application itself. Use the environment-variable script only if other tooling on the device requires the CA through those variables.

Troubleshooting MDM deployment

The following table describes common issues you might encounter when you deploy the application with an MDM solution, and how to resolve them.

The Windows silent installation does nothing

Confirm that the install command uses the /S flag (a capital S), not /silent.

A device reports that the Win32 app is not supported

The device runs Windows Home, which does not support Win32 app deployment. Use a device that runs Windows Pro, Enterprise, or Education.

The IntuneManagementExtension folder or logs are missing

The Intune Management Extension bootstraps only after you assign a Win32 application and the device checks in. Assign the application, sync the device, and then wait a few minutes.

macOS reports a failure although the files installed

If macOS reports installation failure code 0x87D13BA2, the nested helper bundles were left in the **Included apps** list. Delete every helper row so that only the main application bundle (com.aws.QuickWork.mac) remains, and then reassign the application.

Microsoft Intune reports the macOS app as installed but the app is missing

A prior installation left an application bundle or a package receipt on the device. With version-independent detection, this reads as installed and the deployment is skipped. Test on a clean device. To fully clean a test device, remove the application bundle and forget the package receipt, and then test the deployment again.

```
sudo rm -rf "/Applications/Amazon Quick.app"  
sudo pkgutil --forget com.amazon.QuickWork.mac
```

Application binaries are removed after a restart, but profile data survives

This is caused by a version conflict between the MDM-deployed build and the built-in auto-updater, which happens when the software catalog pins an older version while the auto-updater has already applied a newer one. The installer reads the version mismatch as a broken installation and removes the binaries. To resolve it, set **Ignore app version** to **Yes** in the Intune application configuration, and do not pin a specific version while the auto-updater is active. Admin-controlled version pinning is not available at general availability.

A Windows Server or pooled virtual machine does not enroll

MDM enrollment requires a persistent Windows client edition (Windows 10 or 11). Windows Server is not supported for MDM enrollment.

Getting started

This section walks you through downloading, installing, and signing in to the Amazon Quick desktop application for macOS and Windows.

Prerequisites

Verify that your system meets the following requirements.

macOS

Requirement	Minimum
Operating system	macOS 12 (Monterey) or later
Processor	Apple Silicon (M1 or later)
Memory	8 GB RAM
Disk space	2 GB available for installation. 10 GB or more recommended for search indexing and knowledge graph features.
Internet	Required for sign-in, AI model access, and connected services

Windows

Requirement	Minimum
Operating system	Windows 10 (64-bit) or later
Processor	x86_64 compatible
Memory	8 GB RAM
Disk space	2 GB available for installation. 10 GB or more recommended for search indexing and knowledge graph features.
Internet	Required for sign-in, AI model access, and connected services

Downloading and installing

Important

If your organization uses an Enterprise account, complete the steps in [Setting up Amazon Quick on desktop for enterprise deployments](#) before downloading or distributing the application. Enterprise sign-in becomes available after an administrator configures it.

You can download the Amazon Quick desktop application from the following locations:

- **Amazon Quick web application** – Sign in to Amazon Quick on the web. In the left navigation, choose **Extensions**, find the desktop extension card, and choose the download for your platform.
- **Direct download** – Use the following links:

Platform	Download
macOS (Apple Silicon)	Amazon-Quick.dmg
Windows (x64)	Amazon-Quick.exe

To install on macOS

1. Download the .dmg installer file.
2. Open the .dmg file.
3. Drag **Amazon Quick** to the **Applications** folder.
4. Open **Amazon Quick** from the Applications folder or Launchpad.

Note

On first launch, macOS might report that the application is from an unidentified developer or was downloaded from the internet. Choose **Open** to continue. If macOS blocks the application, open **System Settings**, choose **Privacy & Security**, scroll to **Security**, and choose **Open Anyway**.

To install on Windows

1. Download the .exe installer file.
2. Run the installer and follow the on-screen prompts.
3. After installation, launch **Amazon Quick** from the Start menu or desktop shortcut.

Note

On first launch, Windows might display a SmartScreen prompt. Choose **More info**, then **Run anyway**, to continue.

Signing in

When you open the application, the sign-in screen shows the available options.

To sign in with a personal account

1. Open the application.
2. On the sign-in screen, choose **Continue with**, and select your provider (for example, email, Amazon, Apple, Google, or GitHub).
3. Enter your credentials on the provider's sign-in page.
4. Complete any multi-factor authentication prompts.
5. After authentication succeeds, the application loads the Home screen.

Enterprise sign-in becomes available after your organization's administrator configures it.

To sign in with an enterprise account

1. Open the application.
2. On the sign-in screen, choose **Continue with single sign-on (SSO)**.
3. (Optional) Select your AWS Region, or choose **Dynamic** to detect it automatically.
4. Enter your corporate credentials on your organization's identity provider page.
5. After authentication succeeds, the application loads the Home screen.

First-run setup

After you sign in for the first time, Quick guides you through connecting your data sources. The onboarding flow presents categories of connections:

Category	What it connects	Examples
Email	Calendar and email, for meeting prep and action items	A calendar or email connector
Messaging	Threads, direct messages, and channels	A messaging connector
Local files	Folders on your computer	Any folder on your machine

For each category, choose **Connect** to set it up, or **Dismiss** to skip it. You can connect data sources later from **Customize**, on the **Connectors** tab. For more information, see [the section called "Connectors"](#).

Understanding the desktop interface

The interface is organized into a left navigation rail, a main content area, a top bar, and overlay panels.

Home

When you open Quick, the Home screen shows a time-aware greeting and the chat composer as the main entry point. It also surfaces connection setup cards for data sources you have not connected, and Quick Briefs, a summary of your highest-priority feed items with a link to open the full feed.

From the composer you choose which agent handles the conversation, set the response mode (Fast, Balanced, or Smart), and set the thinking effort (Off, Low, Medium, or High). You can also start an Incognito chat, a private conversation that Quick does not save to your history.

Features

The navigation rail on the left provides access to the main areas of Amazon Quick.

Features	What you do there
Chat	Start and continue conversations with Quick. This is the primary way you work.
Feed	Review a prioritized activity feed, your calendar, and your to-dos, drawn from connected services.
Library	Browse, search, and reuse your content and Quick resources in one place.
Customize	Manage the agents, skills, connectors, and knowledge that power your agents.
Apps	Open and build interactive applications (Apps in Amazon Quick).
Spaces	Open curated collections your team maintains, and chat against them.
More	Reach additional areas, such as Research, Data, Dashboards, and Automations.

Choose **Edit navigation** to select and reorder the areas shown in the rail.

Recents

Below the navigation, the Recents area lists your previous conversations, organized into folders and by date. From Recents you can organize conversations into folders, each showing a count; pin a conversation; rename, move, archive, or delete a conversation from its menu; search across your conversations; and use Auto-clean up to sort chats into folders by topic.

Account menu

Selecting your profile at the bottom of the rail opens the account menu, which reaches My context, Settings, Help, and Sign out. Your profile is also where you view your current session usage.

Overlay panels and top bar

The overlay panels and top bar give you access to content and tooling without leaving your current chat view: session activity, the Feed, Mission Control, the All data and apps panel (the spaces, local

folders, and skills available to the current conversation), and Session tabs (an in-conversation reader where documents open beside the chat). The top bar contains a fixed ingress to Mission Control and a connection status indicator.

Working with chat

Chat is the primary way you interact with Amazon Quick on desktop. You can ask questions, get help with tasks, analyze files, create documents, generate visualizations, and work with your connected services through natural language. The following sections describe the chat features available in the desktop application.

Response preferences

Amazon Quick on desktop provides multiple response modes that you choose based on your task. To switch between modes, use the response preferences selector in the chat input area.

Mode	Description	Best for
Fast	Optimized for speed on everyday tasks. Responses are generated quickly with the lowest latency.	Quick questions and simple tasks where speed matters more than depth.
Balanced	Balances quality and speed. Provides high-quality responses with moderate latency.	Most tasks, including writing assistance, summarization, and general day-to-day work.
Smart	Highest-quality responses. Provides the most thorough reasoning, with longer response times.	Complex analysis, nuanced writing, detailed reasoning, and tasks that require high accuracy.

Thinking effort

You control the thinking effort independently from the response mode. Thinking effort determines how much reasoning Quick applies before it responds. Higher thinking effort enables deeper reasoning for complex code, strategic planning, and analysis.

Level	Description
Off	No extended reasoning. Quick responds directly.
Low	Minimal reasoning before responding.
Medium	Moderate reasoning. A good balance for most tasks.
High	Extended reasoning for complex problems.

You can set thinking effort on Balanced and Smart modes. Fast mode responds without extended reasoning. To change the thinking effort, choose the response preferences selector in the chat input area, and then adjust the thinking effort picker.

Chat input

You enter your messages in the chat input area at the bottom of the conversation. From the chat input area, you can attach files, add local folders, turn on web search, attach a space, run commands, and mention agents and items. To open the menu of context options, choose the + (Add context) icon.

Attachments and file uploads

You can attach files directly to your chat messages for Quick to analyze, summarize, or reference. The desktop application supports common file types, including documents, spreadsheets, presentations, images, PDF files, and code files.

To attach a file, use one of the following methods.

- Choose the + (Add context) icon in the chat input area, and then choose Upload files.
- Drag and drop a file directly into the chat window.

Each attachment can be up to 50 MB. When you attach a file, Quick can read, analyze, and reference the file contents in its responses.

Local folders

You can give the agent access to local folders for it to read, write, or search their contents. To add a folder or view the folders you have already granted access to, choose the + (Add context)

icon in the chat input area, choose Local folders, and then choose Manage folders. This opens the Knowledge tab in Customize, where adding a folder works the same as choosing **Customize > Knowledge** and adding a folder there. For more information, see [the section called "Knowledge"](#).

Web search

You can turn on web search to allow Quick to search the internet for up-to-date information when it responds to your messages. When web search is on, Quick can retrieve current information from the web to supplement its responses.

To turn web search on or off, choose the + (Add context) icon in the chat input area, and then choose Web search. When web search is on, Quick includes citations from the web sources it references in its responses.

Spaces

You can attach a space to your conversation to give Quick access to shared team knowledge. Spaces contain curated collections of information that your team maintains in Amazon Quick. While it is attached, Amazon Quick answers from the content in that space and does not reach outside it on its own, so the conversation works within a narrow, scoped context for confidentiality or data compliance. This applies to everything in the space, including its files, datasets, dashboards, topics, and actions.

This is useful when you need answers grounded in your team's shared context, such as the following.

- Project documentation
- Company policies and procedures
- Domain-specific knowledge bases
- Team wikis and reference materials

If the answer is not in the attached space, Quick tells you so rather than answering from somewhere else. Ask Quick to search more broadly and it will, for the rest of that conversation until you change it again. A space you attach applies to that conversation only. Starting a new conversation starts with nothing attached, so you never inherit context from work you have already finished.

To attach a space, choose the + (Add context) icon in the chat input area, choose Spaces, and then select the space you want to use.

Slash commands

You can run commands directly from the chat input area. To open the command menu, type a forward slash (/) at the start of your message, and then choose a command or continue typing to filter the list.

Command	Description
/clear	Start a new conversation with fresh context. Your current chat is saved separately, not lost.
/compact	Condense the current conversation to free up context.
/export	Export the conversation as Markdown or PDF.
/mode	Switch the response mode (Fast, Balanced, Smart) and thinking level (Off, Low, Medium, High).
/rename	Rename the current conversation.
/copy	Copy the most recent response.
/skills	View the available skills.
/connectors	View and manage your connectors.
/goal	Give the agent a job to finish rather than a single question to answer. In a normal chat, it replies once and waits for you. With /goal, it keeps working on its own until the job is complete, double-checks its own work, and only then hands it back.
/help	View the list of available commands.

Mentions

To reference a specific agent, an item in your message, or people from your knowledge graph, type an at sign (@), and then choose from the list that appears. Mentions let you direct your request to a particular agent or person, or point Quick to a specific item, without leaving the chat input area.

Citations and sources

When Quick uses information that is not part of model knowledge to answer a question, it cites where that information came from. Citations help you verify a response and explore the underlying sources.

- **Inline citations:** Quick adds citation markers within its response. Choose a citation to open its source, such as a web page, a file, context, or a connected application.
- **Sources:** Quick shows a Sources summary with the full set of sources that informed a response. Sources can include web pages, files, connected data, and the memories that shaped the response.

For responses informed by your memory, the Sources panel shows the memories that were used. You can boost the memories that were helpful or remove the ones that were not.

Managing messages

You can revise and rerun the messages in a conversation.

- **Edit and rerun:** edit one of your earlier messages and submit it again. Quick generates a new response from the edited message.
- **Retry:** ask Quick to generate a new response to your most recent message.
- **Branch navigation:** when you edit a message or retry a response, Quick keeps each version as a separate branch. Use the branch controls to move between the previous and next versions and compare responses.

Each message also supports inline actions. On a message you sent, you can copy it, edit it, or reply in a thread. Messages can carry emoji reactions, shown as a row of reaction pills. In a read-only transcript, copy remains available while edit and reply are hidden, and reply is not offered inside a thread.

Running work in the background

While Quick is responding, you can send the task to the background and keep working in another conversation. To send the current task to the background, choose the send-to-background control, or press Ctrl+B (Windows) or Cmd+B (macOS). The task continues to run, and Quick notifies you when it finishes.

You can also queue follow-up messages while Quick is working. Queued items show a count and can be sent immediately or removed.

Coding agent sessions

You can use coding agents such as Kiro, Claude, or other connected agents to accomplish coding tasks directly inside a conversation. The desktop application's local code execution runs in an isolated sandbox that blocks common coding commands, so you delegate those tasks to a coding agent. For information about connecting a coding agent, see [the section called “Connectors”](#).

Plan mode

For complex, multi-step tasks, you can use plan mode. When you turn on plan mode, Quick breaks your request into a structured plan with individual steps before it executes. You can review and modify the plan before Quick proceeds with each step. The plan appears in a panel next to your conversation and as a plan card in the chat.

Plan mode is useful for tasks such as the following.

- Multi-file document generation
- Complex data analysis workflows
- Multi-step research projects
- Large-scale code refactoring

Deep research

Deep research provides structured, multi-step investigation. When you start deep research, Quick investigates your question across multiple sources and compiles the results into a comprehensive deliverable. A research plan card appears in the conversation and shows live progress as Quick gathers information, generates figures, and finalizes the report.

Deep research is useful for the following tasks.

- Market research and competitive analysis
- Technical topic deep dives
- Strategic planning and investigation
- Multi-angle problem exploration

To start deep research, ask Quick to research a topic or perform a comprehensive investigation. To follow along while Quick works, choose View research progress on the research plan card.

Incognito mode

Use incognito mode when you want a private conversation that does not update your memory or history. In an incognito conversation, Quick does not read from or write to your long-term memory or knowledge graph, and your responses are not personalized from past conversations. The conversation does not appear in Recents, and it is discarded with nothing retained when you leave it.

Inside an incognito conversation, Quick still follows the current conversation from turn to turn, and it can still search files that you have already indexed. It does not remember anything after you leave the conversation.

To start an incognito conversation, choose the Incognito control at the top right of the chat surface. A private-mode indicator appears while incognito is active.

Note

Privacy is locked once a chat has messages. Turn on incognito before you begin. When you leave an incognito conversation, its contents are deleted and nothing is retained.

Managing conversation context

The Context used indicator shows how much of the conversation context window is in use. As the conversation grows, it fills. To free context during a long session, choose Summarize to compact the conversation history, or start a new conversation.

Away From Keyboard (AFK) recap

When you step away during a long-running task and return, Quick shows an away-from-keyboard (AFK) recap. The recap summarizes what happened while you were away, so you can catch up quickly without scrolling through the full conversation. This recap is specific to a task running in the current conversation. For a summary of activity across your connected services while you were away, see the feed catch-up in [the section called “Activity feed”](#).

Managing conversations

Quick keeps your conversation history in your Amazon Quick account, so it is available to you across the web, mobile, and desktop. You can organize, search, and manage your conversations with the following features.

- **Folders:** organize your conversations into folders. You can create, rename, and delete folders, and drag conversations between them. Auto-clean up sorts chats into folders by topic.
- **Pinning:** pin important conversations to keep them at the top of your chat history.
- **Search:** search across all of your conversations from the search control in the sidebar.
- **Archive:** archive conversations that you no longer need without deleting them.
- **Thread routing:** Quick can route related discussions into organized threads for structured conversations, and you can reply in a thread to branch a sub-conversation from a specific message.

Your conversation history appears in the Recents section of the sidebar, organized by date.

Creating documents and visuals

Quick generates documents directly from chat, including presentations (.pptx), Word documents (.docx), spreadsheets (.xlsx), PDF files, Markdown, and HTML, along with interactive data visualizations. Describe what you need. For example, "Create a presentation about our quarterly results." Generated items appear as artifacts that you can save locally or to your Library. For detailed information about supported formats, editing, visualizations, and best practices, see [Document and visual creation](#).

Quick can also generate and edit images from a text description. Generated images appear as artifacts you can save locally or add to a space.

Session tabs

Documents and artifacts open in session tabs alongside the conversation. In a session tab you view and edit a rendered document, turn on selection mode to select text, images, and tables and add them as context for your next prompt, and act on the open document (open, download, add to a space, share, or open in chat). Multiple documents can stay open as tabs at once.

All data and apps

The All data and apps panel shows the spaces, local folders, and skills available to the current conversation, so you can add or adjust the resources the agent draws on. Skills show a status indicator.

Exporting a conversation

You can export an entire conversation to share it or keep a record. To export the current conversation, use the `/export` command or the export action, and then choose Markdown or PDF.

Activity feed

The activity feed is your single place to catch what happened across your connected services, such as Slack, Outlook, and your calendar, without opening each app. Amazon Quick reads that activity, enriches it with your knowledge graph (the people, projects, and topics in your work), and surfaces only what is relevant to you. A feed agent processes items from your connected messaging, email, and calendar sources, assigns importance, and suggests actions. The feed is organized into Feed, Calendar, and To-dos tabs.

The feed has these characteristics.

- **Personalized every time:** items are ranked by what matters to you, not by what arrived most recently.
- **Burns down as you work:** the feed gets smaller as you handle items, not bigger. The goal is an empty feed, not an endless timeline.
- **Does not consume agent hours:** using the feed does not consume agent hours. Only scheduled tasks that post results to your feed use agent hours. For plan availability and pricing, see [Amazon Quick pricing](#).

On Home

Your activity feed surfaces directly on the desktop and mobile home page, so the items that need you are the first thing you see. From Home you can open the full feed to review everything in one place.

Feed

The Feed tab shows an aggregated activity stream. Each item shows its source service, an AI-generated summary, an importance level, and suggested actions, including a suggested draft when a reply is appropriate. Related items group into collapsible threads. A density control switches the cards between a comfortable and a compact layout.

Quick assigns one of the following importance levels.

Level	Description
Important	Requires immediate attention: direct mentions, urgent requests, time-sensitive updates.
Informational	Relevant but not immediately actionable: project updates, team announcements.
Low priority	Reviewable later: general channel activity and for-your-information messages.

Briefing

Quick on desktop delivers a daily briefing at the top of your activity feed to help you stay on top of your day. The briefing is a curated card that summarizes what you need to know: overnight activity in your connected services, your day plan and calendar, decisions that need your attention, and follow-ups that are due. Choose **Show briefing** on the card to expand it, or choose the card to open the full briefing view.

- **Delivered through the feed:** the briefing appears as a card at the top of your activity feed. You do not need to configure it separately.
- **Updated three times a day:** the briefing refreshes in the morning, midday, and afternoon, so it reflects the latest activity, not just what happened overnight.
- **Personalized over time:** Quick tailors the briefing based on the topics, people, and projects you engage with.
- **Actionable:** suggested actions on the briefing, for example review a draft response, open a document, or start a chat, are the same actions available on individual feed items.

A condensed view of your highest-priority feed items also appears on the chat Home screen, with a link to open the full feed.

Filters and search

Filters are visible by default at the top of the feed. Narrow the feed by label, project, person, or priority, and customize which filters appear for your workflow. A search bar at the top of the feed searches up to seven days of feed activity, so you can locate a specific notification or agent result in seconds. You can also sort the feed, for example ranked, most recent, by project, or by person.

Auto resolve

Quick automatically clears feed items that no longer need you, so your feed stays focused on what is still open. When an item has been handled, such as when you replied to the message elsewhere, the meeting passed, or a newer update supersedes it, Quick resolves it for you. Resolved items collapse behind a Resolved filter, so you can still review what was cleared. You can also give feedback on any item (Not relevant, Already handled, or Something else) to sharpen what surfaces over time.

Catch-up

When you return after being away, the feed offers a catch-up summary: a consolidated view of everything that happened while you were gone, organized by importance level, so you can get up to speed quickly. This summarizes activity across your connected services. For a recap of a single long-running task in a conversation, see the Away From Keyboard (AFK) recap in [the section called “Working with chat”](#).

Always on

Because Quick runs in the cloud, your feed continues to update while the desktop application is closed and your computer is off. Scheduled tasks, monitoring agents, and background workflows run overnight and post their results to your feed, so anything that happened while you were away is captured and waiting for you. You open your day to a feed and briefing that already reflect what is most important, instead of starting from an empty screen or a backlog to sort through.

Calendar

The Calendar tab shows your day on a time grid, today or tomorrow, alongside a Quick Suggests side panel of AI recommendations tied to your events and free time. Each suggestion is a card

you act on, copy, or mark done, and you can start a chat about it. For a free slot, Quick can block focus time directly on your calendar. Other suggestions start a chat with a prefilled prompt, such as preparing for a meeting.

To-dos

The To-dos tab tracks a personal task list alongside the feed. Tasks show due dates and importance, can originate from feed items, and you can create a new task directly.

Configuring the feed

Choose which connected sources surface items to your feed, grouped by function.

Category	Sources
Messaging	Direct messages, mentions, and thread replies from a connected messaging service.
Mail	New email from people who matter to you.
Calendar	Meeting reminders with join links and the day ahead.

You also set how often the feed agent checks for new activity. For feed settings, see Settings.

Library

Library is where you browse, search, and pick up where you left off on content and resources. You can search, filter it by content type, and switch between a table view and a thumbnail grid.

What Library holds

Library stores typed content, including documents, presentations, spreadsheets, PDF files, HTML, images, audio, video, and Markdown, along with Amazon Quick resources such as apps, spaces, agents, skills, research, datasets, dashboards, analyses, flows, automations, and topics.

Views and filters

Switch between list and grid views, with columns for type, owner, and last viewed. Filter by ownership, favorites, owned by you, or shared with you, and by content type. You can favorite items with a star.

Item actions

From an item's menu you open, download, add to a space, share, open in chat, or delete it. Opening an item shows it in a session tab, where you read it and select content to add as context.

Sharing files

Sharing gives view access to that file only. Recipients find the file in their own Library. Choose **Copy link** to get a link to the file that you can send in chat or email. The link works only for people you have already shared the file with.

Apps in Amazon Quick

Apps in Amazon Quick are interactive applications you build by describing what you want, and they are available directly in the desktop application. This section covers how to reach apps on desktop. For how to build, edit, publish, and share an app, see [Build web applications with apps in Quick](#).

There are three ways to get to apps on desktop.

- **From the left navigation:** choose **Apps**, to open the apps landing page, where you describe a new app or browse existing ones.
- **From the Library:** open your Library to see apps you already have access to, including apps shared with you, alongside your other saved items.
- **From chat:** ask Quick to build an app right in a conversation. Quick creates it and opens it for you, with no separate setup needed.

Once an app is open, you can interact with it in a session tab, including full-screen and side-by-side with chat.

Apps in Amazon Quick is available in a subset of AWS Regions. For the list, see [the section called "Regional availability"](#).

Spaces

Spaces are curated collections of documents and Amazon Quick resources that an agent uses as a source of information. On desktop you browse spaces, chat with them as context, and create new ones by navigating to **Spaces** in the left navigation or through chat.

Browsing spaces

Search spaces by name, filter them (Favorites, Owned by me, and Created by me), and switch between Card and List views. Each space offers a quick action to start a chat with it, plus a per-space menu.

Chatting with a space

Attach a space to a conversation to ground the responses from Quick in that shared context, for example project documentation, policies and procedures, or a team reference. You attach a space from the chat composer's data controls or from the All data and apps panel. To learn more about the chat boundary with spaces, see [the section called "Working with chat"](#).

Creating and organizing spaces

Choose **Create** to make a new space. Selecting a space opens it as a tab, so you keep chatting while viewing it. After you create a space, you can create a folder, upload files, upload a folder, or add assets (dashboards, knowledge bases, actions, or datasets).

Uploading your local files and folders makes them searchable in natural language. After your files are uploaded and reach a Ready status, you can query them in natural language.

Note

Adding documents to spaces consumes index cost. For details, see [Amazon Quick pricing](#).

Customize

Customize is the hub for managing everything that powers your agents. It has four tabs:

- **Agents:** build and manage agents. For more information, see [the section called "Agents"](#).
- **Skills:** create and manage skills. For more information, see [the section called "Skills"](#).

- **Connectors:** manage web and cloud connectors, MCP servers, and coding agents. For more information, see [the section called "Connectors"](#).
- **Knowledge:** give Quick access to folders on your computer and control how they are indexed. For more information, see [the section called "Knowledge"](#).

Across the Customize tabs you install additional agents, skills, and connectors from the Add to Quick catalog, which offers items shared with you and, for skills, an Official set published by Amazon that you filter by category.

Agents

Agents are configurable AIs in Amazon Quick. You give each one its own instructions, knowledge sources, and tools so it can answer questions and get work done for you. Quick is the built-in agent, and you can also build your own or install ones shared with you. Each agent can work with you in chat or run on its own on a routine. Agents can only access the same resources and services that you have access to, and you control what each agent is allowed to do.

Users can browse agents that are shared with them on the Agents page by choosing **Browse all**. This view presents available agents in a read-only detail view along with the description and knowledge sources each one uses. From this view, you can install an agent to add it to your agent list. Once installed, an agent is ready to use anywhere you work: you can chat with it, @ mention it inline in any conversation, and run it on a routine, and it appears in your agent switcher for quick access. From the Agents page, you can favorite an agent for easy access or set an agent as your default. A published agent can also be accessed from Quick on the web.

Creating an agent

You can create an agent through the interface or through natural language. To create an agent through natural language, enter "I want to create an agent" in the conversation.

To create an agent

1. Choose **Customize** in the navigation, and then choose the **Agents** tab.
2. Choose **Create**, and then choose **From chat** (or say "create an agent" in chat). To import an existing agent definition, choose **From file** instead.
3. Describe what you want the agent to do in natural language.
4. Quick generates a draft configuration including name, description, and instructions.

5. Review and refine the configuration through conversation or by editing fields directly.
6. Publish the agent when you are ready. Agents cannot be shared with other users until published. Agents are private unless shared explicitly.

You can attach reference documents during creation to provide additional context. You can also start editing an agent by mentioning it with an at-sign in chat, which opens the agent editor in a session tab.

Agent configuration

The agent editor organizes configuration into two tabs.

Instructions tab: set the agent's Name, Description, and Instructions. Instructions support rich formatting (headings, lists, emphasis, code, and quotes) and a welcome message with starter prompts. Instructions can be up to 50,000 characters.

The Capabilities tab has these fields.

Capability	Description
Web search	Toggle to allow the agent to search the web.
Skills	The skills the agent can use. By default an agent can use all skills; if you add specific skills, the agent is limited to those. Remove them to re-enable all skills.
Spaces	Connected data sources the agent reads from.
Connections	Connectors and the tools that the agent can be configured to use.

Using an agent

Use one of the following ways to use your agent.

- **Agent switcher:** choose an agent at the top of chat to start a conversation with it.
- **At-mention:** mention an agent by name in any chat to invoke it inline.
- **Chat button:** start a new conversation from an agent's landing page.

- **Run it on a routine:** pair the agent with a routine to run on a recurring basis without prompting. Routines run in the cloud, so they keep running even when your laptop is closed or offline. The agent stores deliverables in its **Agent Files** and surfaces results on its **Agent briefing** page (when briefing is turned on) and **Mission Control** activity page.

You ask questions and get responses shaped by the agent's instructions and connected data sources, applications, and skills. Have it perform tasks using its configured tools and connections.

Monitoring agents

You can monitor agent activity in two places.

- **Agent landing page:** the agent's overall page, with Overview (including an AI-generated agent briefing), Activities, Agent Files, and Settings tabs. View a specific agent's runs, approvals, routines, and deliverables.
- **Mission Control:** a centralized command center that gives you visibility across all your agents in one place. Monitor all agent activity, unblock agents that need input, view running agents, manage routines, and track KPIs, without switching between individual agent pages. For more information, see [the section called "Mission Control"](#). Access Mission Control from the icon in the top-right navigation bar.

Agent landing page

Each agent has a landing page with four tabs.

Tab	Description
Overview	Summary of the agent's purpose, items needing your input, and an AI-generated agent briefing.
Activities	Activity log with filter tabs: All, Needs input, Running, and Routines.
Agent Files	Persistent file storage: artifacts and deliverables the agent has produced across runs.
Settings	Agent preferences, such as automatic briefing updates. Instructions and capabilities are edited from Configuration, not here.

The Overview tab includes the following sections.

- **Needs input:** when the agent pauses on an action that requires your decision, an approval card appears. Each card shows a description of the proposed action with **Approve** and **Deny** buttons.
- **Agent briefing:** an AI-generated summary of the agent's work (for example, top findings, trend analysis, or status dashboards). It updates after each run and can be regenerated on demand.

The Activities tab shows past and current agent work. You can filter the view using the following sub-tabs.

- **All:** shows all activities in reverse chronological order. Each entry includes a title, summary, agent name, and timestamp.
- **Needs input:** items where the agent paused and is waiting for your decision. Each item shows **Approve** and **Deny** buttons.
- **Running:** tasks in progress.
- **Routines:** lists the agent's configured routines. Time-based routines (schedules), for example, "Every Monday at 9 AM" or "Every weekday at 8:30 AM," are supported on the desktop today. Each routine has a play button to trigger a run and a toggle to turn it on or off. Choose **Create** to add a new routine.

Activities can include file attachments (for example, drafted documents or reports) that you can open and review.

Run status

Each run shows its current status. A run is **Running** while it is in progress, and **Needs input** when it has finished or paused with one or more actions awaiting your decision. Once a run ends, it is **Completed** if it finished with no issues, **Interrupted** if it was stopped before completing, or **Failed** if it ended with an error. An agent that has not run yet shows **No runs yet**. On the Agents list, the Status column summarizes each agent's most recent run, so these finished states appear there as "Last run completed," "Last run interrupted," or "Last run failed."

Managing agents

Agents are versioned. You work on a draft, add release notes that are shown to your team on publish, and then publish. You can manage versions and discard a draft. Agents cannot be shared with other users until published, and a routine always runs the published version.

From an agent's menu you can **make it your default**, so Quick uses it for every new chat until you pick another, as well as **share** it, **export** its configuration to a JSON file, **duplicate** it, **manage its versions** (switch between the published and draft versions), **manage its routines** (the routines that run it on its own), **uninstall** it, or **delete** it. Uninstall removes the agent from your list only; it stays available to everyone else it is shared with. Delete removes an agent you own for everyone. The Agents list supports search, filtering (including by status and version), sorting (including by last run and last modified), favoriting, and list or grid views. The list shows each agent's description, last-modified time, and run status.

For others to use your agent's connected sources, you must share those with the intended recipients. Otherwise, they will only get responses from the sources they have access to. You do not need to share all underlying resources.

Skills

Skills are reusable sets of instructions that teach your agents how to do a specific task. Each skill packages the steps and know-how for that task. Instead of re-explaining how you want a task done every time, you capture it once as a skill, and Quick applies those instructions whenever the task comes up, for example writing a blog post that follows your brand guidelines, or preparing a weekly status update from your template. You can create your own skills, publish and share them, and install skills that other people in your organization share with you. You manage these from the **Skills** tab (**Customize > Skills**).

Skill definition

A skill is a folder with a `SKILL.md` file at its root. That file holds a description (a short summary of what the skill is for, which Quick uses to recognize when the skill applies to your request) and instructions (the steps and guidance Quick follows when the skill is active, which can define a multi-step process, not just a single prompt).

A skill can also reference tools and bundle reference files.

- **Referenced tools:** the tools the skill relies on. Tools are defined separately, as connectors or built-in capabilities. A skill points to the tools it needs rather than containing them, and Quick makes them available when the skill activates.
- **Reference files:** supporting files the instructions draw on, such as scripts, templates, or examples.

How agents load skills

Agents load skills in two stages.

1. Discovery: agents search skills by name and description to know when a skill is relevant.
2. Activation: when your request matches a skill's description, or you name the skill, the agent reads its full instructions into the conversation and then follows them, including the referenced tools and files.

Owner and viewer

What you can do with a skill you created or installed depends on your role for it.

- Owner: you created the skill or duplicated one into your own. You can edit, publish, share, export, and delete it.
- Viewer: someone shared the skill with you. You can add and use it, and it stays in sync with the owner's published version. You cannot edit, publish, re-share, or delete it. To make a version you can change, duplicate it. Duplicating creates a separate skill that you own.

System skills

Quick includes system skills, pre-installed skills for capabilities such as document creation, presentations, spreadsheets, web browsing, code execution, knowledge graph management, and transcription. They carry a System badge. You can try one and turn it on or off, but you cannot edit, duplicate, export, publish, share, or delete it.

Creating a skill

To create a skill, choose **Customize > Skills**, choose **Create**, and then choose a method.

To create a skill from chat

1. Choose **Create**, and then choose **From chat**.
2. Describe the skill you want. Quick drafts the instructions, selects the tools to reference, and creates the SKILL.md file.
3. Review and edit the skill, and then save it.

To create a skill from a file

1. Choose **Create**, and then choose **From file**.
2. Select a .md or .zip file from your computer.
3. Review and edit the skill, and then save it.

To create a skill from a folder

1. Choose **Create**, and then choose **From folder**.
2. Select a folder that contains a SKILL .md file at its root, along with any reference files the skill uses.
3. Review and edit the skill, and then save it.

Tip

When you finish a multi-step task in chat, Quick might offer to save those steps as a reusable skill.

Using a skill

You can use a skill in one of the following ways.

- By name: mention the skill in your message. For example, enter "use the Web Browser skill to check this page."
- Automatic activation: Quick detects when a skill matches your request and loads it without you asking. For example, when you ask Quick to build a presentation, the Presentations skill loads.
- Try it: on a skill's detail panel, choose Try it to start a conversation with the skill loaded. Use Try it to test a skill you created, imported, or have in draft before you publish it.

The skill detail panel

Choose a skill to open its detail panel. The panel shows the skill's description, its instructions, and its reference files. The available actions depend on whether it is a built-in skill or one of your own, your role for it, and whether it is published.

Action	Description
Try it	Start a conversation with the skill loaded.
Edit	Change the skill's instructions and files. Edits are saved to the draft. Owner only.
Publish	Make the draft the live published version, available when the skill has unpublished changes. Owner only.
Duplicate	Copy the skill as the starting point for a new one that you own.
Add to favorites	Pin the skill under the Favorites filter.
Share	Share the skill with people or groups, or with your entire account. Owner only.
Export	Download the skill as a file. Owner only.
Enable or disable	Turn the skill on or off. A disabled skill stays in your list but is not loaded until you turn it back on.
Uninstall or Delete	For a skill shared with you, Uninstall removes it from your list only. For a skill you own, Delete removes it permanently for everyone it is shared with, not just you.

Draft and published versions

A skill you own has two versions.

- **Draft:** the version you edit. Changes you make are saved to the draft and do not affect how the skill runs until you publish. A skill with unpublished changes shows an Unpublished changes indicator.
- **Published (live):** the live published version, which Quick uses when the skill runs.

To update a skill, edit the draft, choose Try it to test it, and then choose Publish. Publishing replaces what is live, and everyone you shared the skill with moves to the update. Because your

draft stays separate from what is published, a skill already in use keeps running the same way while you work on its next version.

Managing skills

The Skills tab lists every skill you can use.

- Search: find skills by name or description.
- Filter: narrow to Favorites, Owned by me, or Shared with me.
- Sort: order the list, for example by Last modified.
- View: switch between a list view and a grid of cards.
- Status: the Status column shows whether each skill is enabled.
- Version indicator: for a skill you own, the list shows a Draft or Unpublished changes label when the draft is ahead of the published version.

Browsing for more skills

To find skills beyond the ones in your list, choose **Browse all** on the Skills tab. This opens Add to Quick, scoped to skills, with two ways to find them: Official (curated, first-party skills provided by Amazon Quick) and Shared with me (skills that people or teams in your organization have shared with you). Search by name, filter by Owner and Category, and sort the results. Choose a skill to preview its description and details before you add it.

When you add a skill from Official or Shared with me, it enters your list as a reference to the owner's published version, not a copy. It stays in sync and updates automatically as the owner republishes. To make a version you can change, duplicate it.

Note

The Quick official skills span everyday productivity, role-based solutions for teams such as Sales, Marketing, and Legal, and industry depth in areas such as Finance, Healthcare and Life Sciences, and Energy and Utilities. New skills are added over time.

Sharing your skills

Publishing and sharing are separate steps. Publishing creates the live version of a skill but does not distribute it. Sharing is a separate action in which you choose who receives the skill. Publishing a skill does not share it with anyone.

To share a skill you own

1. Open the skill's detail panel.
2. Choose **Share**.
3. Choose who to share it with: specific people or groups, or your entire account with Share with all.

People you share the skill with find it in Add to Quick under Shared with me and add it to their own list. As viewers, their skill stays in sync with your published version and updates automatically as you republish.

Note

Your administrator can require approval before a skill is shared, and can control whether skills can be shared with your entire account. See Administrator controls (the next section).

Administrator controls

Administrators can restrict skills through custom permissions. Custom permissions is a scope-down policy that an administrator applies at the account, role, or user level to limit the capabilities people can use. It restricts capabilities; it does not grant them. The following capabilities apply to skills.

- **Skills:** restricts access to skills, including sharing them. This is the parent capability; denying it restricts the skill's full feature set.
- **Share skills with individuals:** a child of the Skills capability. When denied, users cannot share skills with individuals.
- **Share skills with all:** a child of the Skills capability. When denied, users cannot share skills with their entire account; they can still share with specific people or groups if that is allowed.

- Require approval to share skills: a child of the Skills capability. When enabled, a skill must be approved before it is shared.

For how administrators configure custom permissions profiles and how precedence is resolved across the account, role, and user levels, see [the section called “Custom permissions”](#).

Connectors

Connectors bring cloud connectors, system tools, local connectors, and coding agents together in one place, with a catalog to discover and add new ones. A connector links Amazon Quick to an external application or a native tool so agents can access data and act on your behalf. You manage them from Customize, on the Connectors tab, which supports four capabilities.

- Cloud connectors connect Quick to external services such as Outlook, SharePoint, OneDrive, Slack, Jira, Asana, Airtable, and Salesforce, as well as remotely hosted Model Context Protocol (MCP) servers. Once connected, they are reachable from any Quick device, including desktop, mobile, and web.
- System tools are the capabilities Quick provides out of the box, such as image generation, browser automation, and research. You do not install or sign in to them, but you can turn them on or off. Some run in the cloud; others, such as browser automation, are device specific. For the full list, see [the section called “System tools”](#).
- Local connectors run on your own machine, covering local MCP servers you host yourself. Their tools are available only when that machine is on and connected.
- Coding agents such as Kiro and Claude Code can be connected so you can delegate development work to them from chat. See the Coding agents section.

Discovering and adding connectors

On the Connectors tab, you can discover and add connectors that expand what Amazon Quick can do for you. To see the full catalog, choose **Browse all**. Each connector card shows its name and a short description of what it does.

Adding a cloud connector

Cloud connectors in the catalog are either connectors that have been shared with you or connectors that you can set up with a single click of Install followed by a sign-in. The Suggested for

your area contains cloud connectors that have already been set up and shared with you, along with the most popular cloud connectors for Quick users. When you install one of these connectors, by default Amazon Quick connects to the external service using an Amazon-hosted OAuth application for that service. OAuth applications are the standard way an application, such as Quick, makes calls to an external application, such as Outlook or Salesforce, on your behalf. An end-user sign-in is also required.

Some enterprises prefer to use their own OAuth application rather than the Amazon default. To do this, choose Create and then choose Cloud Connector. This Create option gives additional enterprise admin controls during setup, such as setting a connector's name, description, who it is shared with, and its tool permissions. Connectors that do not have an Amazon default OAuth application are also found in the Create section. Creating a custom connection to a remotely hosted MCP server also uses Create and then Cloud Connector.

Connectors you add appear automatically in the connectors table in the lower half of the Connectors tab.

Adding a local connector

Beyond the connectors in the catalog, you can add a local connector that runs on your own machine, such as a local MCP server. A local connector lets Amazon Quick use tools you host yourself, such as a database, an internal API, or a developer tool, so those tools become available to agents in chat and in scheduled tasks. To add one, choose Create and then choose Create local connector. Because a local connector runs on your machine, its tools are available only when that machine is on and connected. For the MCP server fields, see the MCP servers section.

Adding a coding agent

You can also connect a coding agent, an external AI coding tool such as Kiro or Claude Code, that Amazon Quick can hand development work to. Once a coding agent is connected and enabled, you can delegate tasks to it from chat, and it does the work and reports back. To add one, choose Create and then choose Coding agent. See the Coding agents section.

Managing your installed connectors

The Connectors tab has a table of all your installed connectors. Each connector shows its name and description, its Type (Cloud, Local, or System), its Owner, and its Status. To manage your connectors, you can do the following.

Action	Description
Search	Find a connector by name.
Filter	Narrow the list to Cloud, Local, or System connectors. Additional filters are available from the filter control next to the connector type filters.
Sort	Order the list, for example by Recommended.
Sign in	Complete the authentication flow for a connector that requires it before it can be used.
Enable or disable	Turn an individual connector on or off with its status toggle, without removing it.
Create	Add a new connector (an MCP server or a coding agent) using the Create button.
Browse all	Open Add to Quick to discover and install more connectors.

Each connector also has an actions menu with options that depend on its type. For all connectors, you can manage tool permissions (see the next section) and enable or disable the connector.

Connector type	Actions menu options
Cloud connectors	Try it (see example prompts in chat), About connector (view connector details), Sign out, and Uninstall.
Local connectors	Try it, Edit, and Delete.
System tools	Try it.

Some connectors are shared with you by your organization. A connector's Owner column shows whether you own it or it was shared with you. For a connector that supports more than one workspace, for example Slack, choose Set preferred workspace to pick which workspace Quick uses.

Managing tool permissions

Connectors provide tools that agents can use on your behalf. Because a tool can make changes in an external application, Amazon Quick can pause for a consent step before a tool runs, keeping a human in the loop on the changes an agent makes. Whether this consent step is needed is controlled with tool permissions, which are set at two levels: a connector owner level and a user preference level.

Connector owners, usually the application admins who share connectors with users, control one level of settings, and those settings are enforced for everyone who uses the connector. At the owner level, each tool can be set to one of three options.

Owner setting	Description
Let Users Choose	Hands the decision to each user, who then chooses how the tool behaves for them. This is the default.
Always Ask	The tool always pauses for a consent step before it runs, for every user.
Disable	The tool is turned off and is not available to users.

When a tool is set to Let Users Choose at the connector level, each user can set their own preference for the tools they use. The following options are available at the user level.

User setting	Description
Always Allow	The tool runs without a consent step.
Ask Each Time	The tool pauses for a consent step each time it runs.
Disable	The tool is turned off for you and does not run.

By default, tools are set to Let Users Choose at the connector level. The user-level default then depends on whether the tool performs a write task or a read task.

Task type	Connector-level default	User-level default
Write tasks	Let Users Choose	Ask Each Time (can be changed by the user)
Read tasks	Let Users Choose	Always Allow (can be changed by the user)

New tools discovered during an MCP sync are disabled by default, so a connector owner can review them before making them available.

MCP servers

Amazon Quick supports the Model Context Protocol (MCP), an open standard that extends Quick with custom tools and integrations. When you connect an MCP server, its tools become available in your chat conversations and scheduled runs, for example a server that provides database query tools.

To add one, choose Create and then MCP server. A dialog offers three connection types. You can also paste a standard MCP server JSON configuration to populate the form.

Local

Run an MCP server as a command on your machine, the most common option for development tools and locally installed servers.

Field	Required	Description
Name	Yes	A friendly name, for example "My Database MCP".
Command	Yes	The executable to run. Common values: python, npx, node, uvx.
Arguments	No	Command-line arguments, space-separated, for example <code>-m mcp_server --port 8080</code> .
Description	No	What the server does and what tools it provides, so Quick knows when to use it.

Field	Required	Description
Environment variables	No	Key-value pairs the server requires.
Startup timeout	No	Seconds to wait for the server to start, 5 to 300 (default 30).

Import

Load MCP server definitions from an existing configuration file, so you reuse configurations from other tools. Quick supports configuration files from Kiro, Claude Code, AIM, Antigravity, and QuickWork exports. Enter the configuration file path (for example, `~/.kiro/settings/mcp.json`). Quick scans the file and detects the servers defined in it, and also detects compatible tools installed on your system, showing them as chips under "Detected on this system". Choose Load file to import the servers.

Remote

Connect to an MCP server over HTTP, for shared team servers or hosted services.

Field	Required	Description
Name	Yes	A friendly name.
URL	Yes	The MCP endpoint URL. Must start with <code>http://</code> or <code>https://</code> .
Headers	No	HTTP request headers as key-value rows, for example an <code>Authorization</code> header. A header value can be stored as a secret in your operating system keychain.
Description	No	What the server does and what tools it provides.
Startup timeout	No	Seconds to wait for the server, 5 to 300 (default 30).

After you add servers, they appear in the MCP servers section, where you search, filter by status, toggle individual servers on or off, edit, and remove them. You can attach MCP servers to a schedule so automated runs use their tools.

Coding agents

Coding agents are external AI coding tools that Amazon Quick delegates development tasks to, using the Agent Client Protocol (ACP). When enabled, you ask Quick to send work to a coding agent, for example "refactor this module," and Quick reports the results back.

To add a coding agent

1. Choose Customize, and then choose the Connectors tab.
2. Choose Create, and then choose Coding agent.
3. Configure the agent, and then save.

You add a coding agent from a preset (a one-click add for a supported agent, which you edit afterward) or with a custom configuration. A custom agent is configured with a Name, the launch Command and Arguments, an optional Description, and optional advanced settings.

- Environment variables: values the agent needs at launch.
- Override path: where the agent runs. When empty, the agent runs in the current conversation's workspace. An absolute path supports mounted, external-volume, or remote setups.
- MCP server forwarding: which MCP servers are forwarded to the agent: None (the default), All enabled, or Choose specific servers.

For a supported preset agent, additional configuration is available: tool permissions with three states per tool (Off, Ask for Permission, or Auto-approve), write-path scoping, and model selection. You can test the connection before saving.

Add only coding agents your organization approves. You are responsible for the security and data-handling compliance of any coding agent you register.

Knowledge

The Knowledge tab in Customize is where you give Amazon Quick access to files on your computer and control how Quick indexes them. For each folder, choose whether Quick uploads its files for

full-content search, extracts information from them into your knowledge graph, both, or neither. Adding a folder grants access only; indexing stays off until you turn it on.

Indexing options

Open a folder's details pane and choose **Settings** to find these options under Indexing.

Option	Description
Allow full file context for enhanced searching	Uploads and indexes this folder's files so Quick can search across their full content, including in scheduled and background work. When this option is on, it shows estimated monthly usage as a percentage of the monthly indexing budget.
Always remember file information	Uploads this folder's files and extracts people, projects, and dates from them into your knowledge graph. When this option is on, it shows estimated monthly usage as a percentage of the monthly budget available for extraction.
Agent access	Lets Quick read files from this folder. A folder in your list always has agent access, so turning this off removes the folder. Quick asks you to confirm first.

Details

The Details tab reports Files, Size, and Entities extracted. Entities extracted shows a dash until you turn on Always remember file information. Choose **Files** to jump to the advanced settings. Choose **Sync** at the top of the pane to pick up changes immediately.

Advanced settings

Choose **Advanced** on the Settings tab to reach the per-folder indexing controls. Scan interval (every 30 minutes) and Max file size are read-only; files larger than the maximum are skipped. You can set Max file age to index only files modified within a window, restrict File extensions, add Exclude patterns as glob patterns that Quick never indexes, and give knowledge graph extraction instructions to guide entity extraction. Rescan folder picks up changes incrementally. Rebuild index rebuilds from scratch and shows a cost estimate first.

Search indexing limits

Global search-indexing limits apply across all indexed folders. Indexing stops automatically when free disk space falls below 8.0 GiB. You set a storage limit for the knowledge database, a maximum file size for indexing (larger files are skipped but remain available through direct search), and a maximum folder size for indexing.

Folder access permissions

To set what Quick can do with a folder, choose **Manage permissions** and set an access level: Full Access (read and write), Read Only, or Ask Each Time. You can adjust individual read and write operations independently.

Mission Control

Mission Control is a centralized command center in the Amazon Quick desktop application that gives you full visibility and management over all your agents in one place. You monitor agent activity, review outputs, manage routines, and track performance metrics, so you do not need to check on agents individually.

Why use Mission Control

As you scale your use of agents with background tasks and automated workflows, keeping track of everything becomes complex. Mission Control provides these benefits.

- Centralized visibility: view all agent activity without switching between conversations or agents.
- Faster unblocking: agents that need your input are surfaced clearly, so you respond quickly.
- Operational awareness: key performance indicators (KPIs) and run history confirm your agents are working correctly.
- Efficiency at scale: whether you have 2 agents or 20, Mission Control keeps you in control.

Capabilities

Mission Control provides these capabilities for managing your agents.

Capability	Description
Monitor all agent activity	View a unified list of what every agent has done across all runs.
Filter by agent name	Drill into a specific agent's history and access its outputs.
Unblock agents	Filter on agents that need your input and provide it directly.
View running agents	See which agents are currently active and monitor their progress.
Manage routines	View and edit routines for all agents in one place.
Track KPIs	Understand error rates, active agents, completed tasks, and total runs.
View creator analytics	See insights on how your shared agents and apps are being adopted.

Routines

The desktop application runs your agents automatically through routines. A schedule, an agent that runs at recurring times, is the first routine type. You define what the agent does, when it runs, its capabilities, and its response mode. Quick delivers each run to your feed or a connected app. Schedules are currently the available routine type.

You configure a routine in a modal with two tabs: Overview and Activity.

The Overview tab has these fields.

Field	Description
Repeats	How often the routine runs (for example, daily) and the time of day it runs.
Custom time range	Optionally set a start date and time, an end date and time, and a time zone that bound when the routine is active.
Triggers	Select the agent that you want to run when the routine starts.

Field	Description
Uses	The capabilities, such as MCP servers and tools, the routine can use. Choose + Add to attach more.
Instructions	What the agent should accomplish each time the routine runs.
Model details	Expand to configure the model, thinking effort, memory, and sub-task settings for the routine.

Expand Model details to configure these options.

- Model: choose Fast, Balanced, or Smart.
- Thinking effort: have the agent apply additional reasoning before responding.
- Include memory: allow the agent to use profile and search tools.
- Allow spawning sub-tasks: let the agent break work into parallel sub-tasks.

The Activity tab shows the run history for the routine. From the modal, you can also enable or disable the routine, run it immediately, edit it with chat, and choose Done to save.

Quick includes pre-configured routines, such as a feed agent that processes items from your connected services and populates the activity feed. You customize a built-in routine's cadence, capabilities, instructions, and model.

To create a routine

1. Open Mission Control from the top bar.
2. Choose Create to define a new routine.
3. Configure how often it repeats, its capabilities, instructions, and model.
4. Turn it on when you are ready for it to run.

You can also ask Quick in chat to create a routine. For example, "check my channels every morning at 9 AM and summarize what I missed."

Routines run in the cloud, so they run at their configured time even when the desktop application is closed or your computer is off. The exception is a run that needs a local tool, such as reading a

file on your machine, browser automation, or local code execution. For those steps, your computer must be on with the desktop application running and connected at that time. Otherwise the run proceeds but the local step is unavailable until the desktop reconnects.

Working with agent runs

Mission Control organizes agent activity into runs. Depending on the state of a run, you can take different actions.

For a completed run, you view what the agent did and access the artifacts it produced, get an in-depth summary, track events of all sub-tasks, and chat to follow up on results. For an ongoing run, you view live events as the agent works, pause the run, or cancel it. For a run that needs input, you approve or deny authorization requests for third-party applications, or chat with the agent to provide clarifications.

KPIs and metrics

Mission Control displays key performance indicators that help you understand agent health and productivity at a glance. These metrics are available.

Metric	Description
Failed runs	Number of agent runs that ended with an error.
Active now	Count of agents that are currently running.
Routines completed	Number of routines completed today.
Runs this week	Total agent runs and app visits in the selected period, showing how your deployed agents and apps are being adopted.

Viewing your runs and creator analytics

From the Mission Control landing page, choose the Runs this week KPI to open the activity page. The activity page has two views: Runs and Creator analytics.

The Runs view shows your agent and app run activity for the selected time range (this week, last week, this month, last month, or last 90 days). It displays:

- **Summary cards:** a Runs this week card (a total with a trend chart broken down by Agents and Apps) and an Agent hours card (the amount of agent runtime used, with the reset period).
- **Runs list:** a searchable, sortable list of runs that you can filter by all, agents, or apps. The list has the columns run name, run type (agent or app), uses, run status, and last run (shown as last visited for apps). Rows expand to show more detail.

The Creator analytics view shows activity across the assets you have published or have co-owner access to, for the selected time range (this week, last week, this month, last month, or past 12 months). You can filter by all, agents, or apps.

- **All:** summary cards for total assets shared and total assets used (each broken down by agents and apps), and a table with the columns asset name, type, usage, active users, and shared with.
- **Agents:** summary cards for total agents shared and total conversations, and a table with the columns agent name, total active users, total conversations, average queries per user, and shared with.
- **Apps:** summary cards for total apps shared and total app views, and a table with the columns app name, total active users, total app views, and shared with.

A Last updated timestamp shows when the analytics were refreshed. The creator analytics page refreshes every hour.

Accessing Mission Control

You can access Mission Control from anywhere in the desktop application. Choose the Mission Control icon in the top-right navigation bar. A quick-access popover appears showing agent runs, pending inputs, and routine management. To view the full Mission Control landing page, choose the Mission Control header or the View all button in the popover. Work from the lightweight popover for quick actions such as unblocking an agent, or open the full page for a comprehensive view of all agent activity.

Examples of routines

- **Morning brief:** every morning, summarize unread messages, new email, and today's calendar.
- **Inbox triage:** every 30 minutes, scan new email and flag anything urgent.
- **Project monitor:** daily, check a channel for updates on a project and compile a summary.
- **Meeting prep:** before each meeting, gather relevant context and prepare a brief.

- **Competitive monitor:** weekly, search the web for news about specified companies and produce a report.

My context

My context is where you review and manage what Amazon Quick knows about you and your work. You reach it from the account menu. It has two tabs: [the section called “Memories”](#) and [the section called “Knowledge graph”](#).

Memories

Memories are the facts, preferences, procedures, and patterns Amazon Quick learns from your conversations to personalize future responses. Each memory shows a confidence score, how many times it has been referenced, and its source.

Memories are kept per agent. The default assistant and each custom agent you create have their own memories, so a rule you set for one agent does not silently change how another agent behaves. Entities in your knowledge graph remain shared across agents.

Reviewing and organizing memories

Search memories, filter and sort them, and switch between list and grid views. A single Filters panel covers the following.

- **Agent:** use the Agent dropdown at the top of the tab to view memories for All agents, Default (chat), or a specific custom agent.
- **Type:** filter by memory category. The available categories are supplied at runtime, so they vary. Examples include procedures, preferences, and profile facts.
- **Provenance:** show only Told memories (you stated it explicitly to Quick), only Learned memories (Quick inferred it from your activity or conversations), or both.
- **Sort:** Date created, Accuracy confidence, or Times referenced.

Each memory card shows the agent it applies to. Knowledge graph entities in the memory panel are marked shared to indicate that they are visible to every agent.

A memory's detail view shows its text, an information section (created, updated, times referenced, confidence, type, and Applicable to), an attributes section (the Told or Learned badge, any tags,

and a global-behavior badge), and a step-by-step provenance timeline that traces what Quick used to form it. Each memory card provides the following actions.

- **Boost:** raise the confidence in the memory to the highest level so it is prioritized in future responses.
- **Forget:** remove the memory. When you forget a memory, you can optionally choose a reason (for example, Too specific, Never said this, or Outdated) or add free-form feedback. Quick uses the reason to avoid re-learning the same thing.
- **Edit:** update the memory's text. Use edit to correct a fact rather than delete and re-teach it.

Memory settings

The Memory section in Configuration includes a master switch and a Forget all memories action.

- **Enable memory:** a master toggle that controls whether Quick learns from your conversations and personalizes future responses. When memory is off, existing memories are not used and no new memories are recorded.
- **Forget all memories:** permanently deletes every memory Quick has saved. Type "forget" to confirm.

Forgetting all memories cannot be undone. It does not affect your knowledge graph. To clear entities and relationships, use Reset knowledge graph (see [the section called "Knowledge graph"](#)).

Knowledge graph

Amazon Quick builds a personal knowledge graph that captures entities and the relationships between them from your connected data sources and files. The graph gives Quick contextual understanding of the people, projects, organizations, and events in your work, so it provides more relevant, personalized responses. It is kept for you, per user, in your Amazon Quick account.

If you ask about a project, Quick uses the graph to identify related people, recent meetings, associated documents, and relevant channels.

Entity categories

Quick organizes graph entities into categories that span the people, organizations, projects, places, events, communications, documents, actions, decisions, and terminology in your work, along with

structured data types such as datasets, dashboards, and topics when relevant. The category set is supplied at runtime. Each entity has a count of edges, the relationships connecting it to other entities.

Data sources

The graph extracts entities from connected integrations and from local folders. When you enable auto-ingestion, Quick extracts entities from your connected apps. When you enable knowledge graph extraction for a folder on the Knowledge tab in Customize (see [the section called “Knowledge”](#)), Quick extracts entities from files in that folder, independently of keyword and semantic indexing.

Viewing the knowledge graph

Open My context and choose the Knowledge graph tab for an interactive, force-directed visualization where nodes are entities and edges are relationships.

- Node color: each category has a distinct color; a legend maps colors to categories and lets you highlight a category.
- Node size: scaled automatically by how connected an entity is; more-connected entities appear larger.
- Search entities: find an entity by name and jump to it.
- Interact: click a node to select it and view its details, double-click to focus on a node and see its direct connections, scroll to zoom, and drag to pan.
- View mode: switch between Top N connected (the most-connected entities) and Communities (clustered by community detection).
- Display limits: choose how many nodes (50 to 1,000) and edges (100 to 2,000) render, shown against the totals, to keep large graphs responsive.
- Node details: selecting an entity opens an entity details panel (see the next section).

Working with an entity

When you select an entity, the entity details panel shows the following.

- Name and category: the entity's display name and its category (for example, Person or Project). Choose Edit next to either to rename the entity or change its category.

- **AI-generated summary:** a short summary of what Quick knows about the entity. Choose **Regenerate** to rebuild the summary from the latest data.
- **Connected entities:** other entities that this entity is related to, grouped by relationship.
- **Source files:** the files and messages Quick used to extract this entity.
- **Ask:** start a new chat about the entity. Quick opens a conversation preloaded with the entity's context.
- **Delete:** remove the entity and its relationships from the knowledge graph. This permanently removes the entity and any relationships that connect to it, and cannot be undone.

Configuring and resetting the knowledge graph

The Knowledge graph settings include Auto-ingest from integrations, a single toggle that lets Quick learn from your connected apps and build the graph. Per-folder knowledge graph extraction is configured separately, on the Knowledge tab in Customize (see [the section called “Knowledge”](#)).

You can reset the knowledge graph to permanently remove all entities and relationships that Quick has built. Resetting removes all people, projects, organizations, events, and relationships. It does not affect your files, your memories, or your connected services. Quick can rebuild the graph from those sources over time if you leave auto-ingest and per-folder extraction enabled.

To reset the knowledge graph

1. From the account menu, choose **My context**.
2. Choose **Configuration** in the upper-right corner.
3. In the Knowledge graph section, choose **Reset knowledge graph**.
4. Type "reset" to confirm, and then choose **Reset**.

A reset is permanent. There is no undo. Consider narrowing auto-ingest or per-folder extraction first if you only want to remove a specific source of entities.

Settings

Open Settings from the account menu. Settings are organized into tabs: Account, Appearance and language, Chat, General, and Advanced.

Account

Manage account identity and view plan and usage. The Account tab shows your profile, and, for a managed account, an indicator that your organization manages it, a session-usage meter for agent hours, and an Account region selector that controls where account data is stored and processed.

Appearance and language

- Color mode: Light, Dark, or System.
- Language: the interface language.
- Timezone: used for schedules, timestamps, and calendar availability.

Chat

- Instructions: global instructions sent with every message across all agents, such as preferred language, format, and style.
- Default agent: the agent used for new chats.

General

- Desktop notifications: show operating-system notifications when Quick is in the background.
- Browser: choose the default browser Quick opens for web tasks and automation (System default, Chrome, Edge, Brave, or Firefox), and turn on Use my existing browser to connect Quick to your running browser, where your logins, cookies, and extensions are available.
- Keyboard shortcuts: a screenshot-capture shortcut and the key combination that sends a message, including a custom option you record.
- Maximum parallel tasks: a slider from 1 to 20 that controls how many background tasks run at once.
- Updates: install updates automatically, so Quick downloads new versions in the background and installs them on restart.

Use my existing browser

When you connect Quick to your running browser, follow the setup steps to enable remote debugging, then test the connection. Enabling remote debugging allows applications on your

computer to connect to the browser and access open pages, cookies, and authenticated sessions; turn it off when you are not using this feature. When this option is off, the agent launches a separate browser instance with a copy of your profile.

Advanced

- **Export diagnostics:** save diagnostic logs for troubleshooting; choose a time range and export.
- **Clear all data:** removes all Quick data, including conversations, cached messages, knowledge graph, saved credentials, and preferences. The application quits after cleanup. Clearing all data is irreversible.

System tools

Amazon Quick on desktop includes built-in system tools that provide core capabilities such as web search, file operations, browser automation, image generation, and code execution. You enable or disable each tool individually and configure granular permission controls. Manage system tools in **Customize**, on the **Connectors** tab.

Managing permissions

In the connectors table, each system tool has an enable/disable toggle and an actions menu. Choose the ellipsis (actions menu), and then choose **Manage permissions** to set the tool's access level and adjust its per-operation permissions. To turn a tool on or off, use its enable/disable toggle.

Access levels

You set one of three access levels:

Access level	Description
Always Allow	Quick has full read and write access to this tool. Operations run without prompting you for confirmation.
Ask Each Time	Quick asks for your confirmation before performing any operation with this tool.
Always Deny	Quick cannot perform any operation with this tool.

Below the access level, each tool lists its operations grouped into read and write, and you set each independently, for example Always Allow overall but a specific write operation set to Ask Each Time.

To manage permissions for a tool

1. Choose **Customize** in the navigation, and then choose the **Connectors** tab.
2. Find the system tool in the connectors table, choose the ellipsis (actions menu), and then choose **Manage permissions**.
3. Set the access level, and adjust individual operation permissions.
4. Close the panel. Changes take effect immediately.

Available system tools

The built-in system tools include:

- **Web Search:** find current information from the internet through search and fetch-URL operations, with citations.
- **File operations:** download files from URLs and open local files in their default applications.
- **Browser Automation:** launch a browser, navigate, interact with elements, fill forms, take screenshots, and extract text. It supports a default mode, which uses a separate browser instance with a copy of your profile, and a use-your-own-browser mode.
- **Image Generation:** generate and edit images with Amazon Nova Canvas, including remove background, variations, and outpaint (extending an image beyond its original boundaries).
- **Engram Builder:** analyze your writing across messages and conversations to build a personality profile (an engram), so Quick can write in your style. Operations: get, list, save, and delete engrams.
- **Knowledge and memory:** query and build the knowledge graph, and recall learned patterns.
- **Agent management:** create, update, delete, and trigger routines from chat.
- **Task management:** break complex requests into parallel background sub-tasks, organized into groups, with progress tracking.
- **Chat and notifications:** surface contextual suggestion pills, message reactions, desktop notifications for events such as task completion, and briefing summaries on the Home screen.
- **Transcription:** convert audio and video files to text so Quick can summarize, quote, or answer questions about their contents. The Transcribe media operation transcribes an audio or video

file (for example, .mp3, .mp4, .wav, .flac, .m4a, or .ogg) to text. Use it when you want Quick to work with recordings such as meeting captures, voice memos, screen recordings, podcast clips, or interview audio as if they were text.

- **Explainability:** produce a structured account of how Quick arrived at its answer. The `Emit` operation returns an explanation that includes assumptions (what Quick took as given), sources (which datasets, files, connectors, or memories were consulted), calculations (the steps that produced the result), and limitations (known caveats or places where the answer is uncertain). Use it when you want to see the reasoning behind an answer, share it, or check the assumptions before acting.
- **File indexing and search:** find and cite content across the folders you have granted access to. Operations include `Search files` (search across indexed files by keyword or meaning), `Read document` (return the full content of a specific file), `Check indexing` (report indexing status for a folder or file), and `Approve indexing` (approve or skip a pending indexing operation). For how folders are indexed, see [the section called “Knowledge”](#).
- **Working with your data:** answer questions in natural language about the datasets and dashboards you connect through Quick spaces. Quick can answer questions from datasets (run analytical queries and return an answer with the underlying data), read dashboards (interpret and summarize the data behind a dashboard visual), and discover data (browse the catalog of datasets and dashboards available to you). Which datasets and dashboards Quick can see depends on your organization's permissions.

All system tools are enabled by default with Always Allow permissions. You restrict or disable any tool at any time; changes take effect immediately for new conversations and agent runs. For the permission model and hardening guidance, see [the section called “Security and privacy”](#).

Code execution

The Code Execution tool writes and runs Python in a sandbox for data analysis, calculations, file processing, chart generation, and automation. A persistent namespace carries variables across runs within a conversation, and pre-installed packages are available. Sandbox file access follows your folder permissions, plus system temporary directories.

Security, privacy, and architecture

Amazon Quick on desktop is designed to keep your data private while providing full access to AI capabilities.

Your data is never used for AI model training. AWS does not use your conversations, files, or personal context to train or improve AI models.

How data is handled

Amazon Quick runs in the cloud. Your conversations, memory, knowledge graph, agent-produced files (your Library), and the index of content you make searchable are kept per user in your Amazon Quick account, isolated from other users. The desktop application stores only what it needs to operate locally: application configuration (settings, preferences, and connection state), the folders you have granted access to and their local settings, locally configured MCP and coding-agent servers, and your sign-in credentials (authentication tokens, held in your operating system's credential vault).

Data storage

Local application data, configuration, folder permissions, locally configured MCP and coding-agent servers, and sandbox settings, is stored in `~/ .quickwork/` on macOS or `%USERPROFILE% \ .quickwork\` on Windows. Your conversations, memory, and knowledge graph are not stored here; they are kept in your Amazon Quick account.

Folder permissions

Quick uses operating-system sandboxing to control file access. It accesses only folders you explicitly grant, and you revoke access at any time. Each folder supports independent controls for keyword search, semantic search, and knowledge graph extraction, plus granular per-operation permissions.

Quick also has access to system temporary directories regardless of folder permissions. On Windows these are `C:\TEMP`, `C:\TMP`, `\TEMP`, and `\TMP`. On macOS and Linux these are `/tmp`, `/var/tmp`, and `/usr/tmp`.

System tool permissions

Amazon Quick on desktop governs three kinds of capabilities with the same permission model: system tools that provide core capabilities, connectors to third-party services such as Slack, Outlook, and Gmail, and local tools that run on your device. You can turn each capability on or off. Each supports a three-tier permission model, Always Allow, Ask Each Time, or Always Deny, with granular per-operation controls (see [the section called "System tools"](#)).

Quick stores your permission choices and applies them consistently whenever an agent acts on your behalf, including scheduled agents. An agent can only take actions you have authorized, and you can review or change these permissions at any time.

Connection security

- OAuth 2.0: connected services authenticate with OAuth 2.0; Quick never sees or stores your third-party passwords.
- Independent connections: each service is managed independently; you disconnect and reconnect any service without affecting others.
- Minimal permissions: Quick requests only the permissions needed for each service.

Privacy controls

Quick provides controls for whether it learns from your conversations, searches your conversation history, and extracts entities from connected services. You view, edit, and delete individual memories (see [the section called “My context”](#)).

Network access and required domains

The application makes outbound connections for discovery and data-plane operations, remote configuration, application updates, and identity-provider authentication. In restricted networks, add the required Amazon Quick domains to your allow list.

Category	Domain	Wildcard meaning	Purpose
Amazon Quick console	*.quicksight.aws.amazon.com	* = AWS Region (for example, us-east-1)	Web console, resource discovery
Amazon Quick data plane	*.dp.integrations.*.prod.plato.ai.aws.dev	First * = per-account cell ID; second * = AWS Region	Inference, streaming, search

Category	Domain	Wildcard meaning	Purpose
Amazon Quick enterprise gateway	*.desktop .enterprise.quick. aws.dev	* = multi-level subdomain (<cell>.<stage>.<region>)	Enterprise account bidirectional communication
Remote configuration and updates	*.cloudfront.net	* = CloudFront distribution ID	Feature flags, application updates, Apps in Amazon Quick content
Telemetry	cognito-identity.* .amazonaws.com	* = AWS Region	Operational telemetry
Identity provider	Customer-specific (for example, login.microsoftonline.com)	Not applicable	Enterprise SSO authentication

For strict proxy environments that cannot use wildcards, you can determine the account-specific cell hostname after you first sign in. It follows the pattern <cell-id>.dp.appintegrations.<region>.prod.plato.ai.aws.dev, where <cell-id> is a stable per-account identifier and <region> is the account's home region. If your organization inspects encrypted traffic, the operating-system trust store must trust the inspection certificate authority, or the Quick domains must be excluded from inspection (see [the section called "Enterprise setup"](#)).

Clearing all data

To completely reset Quick, use Clear all data in Settings, on the Advanced tab. This is irreversible and removes all conversations, knowledge graph data, saved credentials, and preferences.

Troubleshooting

Installation and launch issues

- macOS reports an unidentified developer: open System Settings, choose Privacy & Security, and choose Open Anyway.
- The application does not launch: restart, verify the minimum requirements, and check disk space. As a last resort, delete ~/ .quickwork/ (macOS) or %USERPROFILE%\ .quickwork\ (Windows) and reinstall. Deleting this directory removes all local data.

Connection issues

- A messaging, email, or calendar connection fails: disconnect and reconnect in Customize, on the Connectors tab, verify your permissions in the third-party service, and check your internet connection.
- A connection shows Not synced: verify connectivity, then disconnect and reconnect; if it persists, the service might be having an outage.
- Sign-in or responses fail only on a corporate network or virtual private network (VPN): confirm the required Quick domains are reachable, and address Transport Layer Security (TLS) inspection with your network administrator (see [the section called “Security and privacy”](#)). If it persists, export logs and contact AWS Support.

Performance issues

- Slow responses: switch to Fast mode, reduce thinking effort to Off or Low, close unused applications, and check your connection.
- High memory usage: reduce the number of indexed folders, disable semantic search where you do not need it, lower Maximum parallel tasks, and restart to clear cache. Search indexing pauses automatically when free disk space falls below 8.0 GiB.

File access issues

- Quick cannot find a file: confirm the parent folder is added on the Knowledge tab in Customize. System temporary directories are always accessible. After you add a folder, Quick reads it immediately; indexing runs in the background.

- Search returns no results: confirm keyword search is on for the folder, check that indexing status is Ready, confirm the file is within the maximum index size, and enable semantic search for natural-language queries.

Browser automation issues

Verify Browser Automation is enabled in Customize, on the Connectors tab. If you use your own browser, complete the remote-debugging setup and test the connection; otherwise, confirm a supported browser is installed.

Schedule issues

- A schedule did not run: confirm the schedule is enabled in Mission Control. Schedules run in the cloud, so they run even when your computer is off; but a run that needs a local tool (file access, browser, or code execution) requires your computer to be on with the desktop running and connected at that time.
- Unexpected results: in Mission Control, review the schedule's instructions and capabilities, consider a higher model (Balanced or Smart), and run it manually to observe results.

Managing agent-hours consumption

- Reduce the number of background schedules; each run consumes agent hours.
- Match the model and thinking effort to task complexity: Fast with thinking Off for simple tasks; Balanced or Smart with higher thinking levels consume more agent hours.

Session usage

View current session usage by selecting your profile. The session limit is a service-protection limit that spreads your monthly agent-hours allocation over time. When you reach it, the application indicates the allocation has been used, and the limit recovers gradually.

MCP server issues

- Local: verify the command and that the executable is installed, check argument formatting, raise the timeout, and verify environment variables.

- Remote: verify the URL and that the server is running, verify any authentication header, and check network and firewall settings.
- Imported: verify the configuration file path and that the source application is installed.

Exporting diagnostics

Export logs from Settings, on the Advanced tab, by selecting a time range and choosing Export Diagnostics. If you cannot open the application, find logs at `~/Library/Application Support/QuickWork/` and `~/Library/Logs/QuickWork/` (macOS) or `%APPDATA%\QuickWork\` and `%APPDATA%\QuickWork\logs\` (Windows).

Getting additional help

Export diagnostics and share them with your support team, contact your organization's IT administrator for enterprise issues, and visit the Amazon Quick documentation for the latest guidance.

Amazon Quick on mobile

Amazon Quick is available as a companion app for iOS and Android that works alongside the desktop application. Use the mobile app to chat with Quick, review your activity feed, and open dashboards while you are away from your computer. You sign in with the same account, and your conversations stay in sync across your devices.

The mobile app focuses on the tasks you are most likely to do on a phone. Some capabilities that are available on desktop are limited or not available on mobile. This chapter describes what you can do on mobile and how the mobile experience differs from desktop.

Topics

- [Getting the app](#)
- [Signing in](#)
- [Feature availability at a glance](#)
- [Mobile-specific behaviors](#)
- [What is not available on mobile](#)
- [Differences between iOS and Android](#)

Getting the app

Download Amazon Quick from the [App Store](#) on iOS or from [Google Play](#) on Android. After you install the app, open it and sign in to get started.

Note

Some capabilities require a recent version of iOS or Android. To get the full mobile experience, keep the app and your operating system up to date.

Signing in

Sign in to the mobile app with your organization's Amazon Quick credentials. The app supports enterprise single sign-on (SSO), including IAM Identity Center federation.

For added security, you can turn on biometric app lock (Face ID, Touch ID, or fingerprint) so that Amazon Quick requires biometric verification each time you open it. You manage this in Settings, under Security. Your administrator can also require biometric lock for your organization.

Feature availability at a glance

The following tables summarize feature availability on mobile compared to desktop. Availability is one of the following.

- Available: works on mobile.
- Limited: partially available on mobile. See the note for details.
- Not available: not available on mobile. Use the desktop application for this capability.

Chat

Feature	Availability
Chatting with Quick	Available
Response modes (Fast, Balanced, Smart)	Available
Thinking effort	Available
Attachments and file uploads	Available
Web search	Available
Mentions (@)	Available
Citations and sources	Available
Editing and retrying messages	Available
Branching conversations	Available
Plan mode	Available
Deep research	Available

Feature	Availability
Managing conversations (folders, pinning, search, archive)	Available
Approvals	Available
Running work in the background	Available
Incognito mode	Not available
Slash commands	Not available
Attaching a space	Not available
Context meter	Not available

Note

Branching conversations is available on Android only. See [the section called “Differences between iOS and Android”](#). For how these chat capabilities work in detail, see [the section called “Working with chat”](#).

Creating documents and visuals

Feature	Availability
Document creation	Available
Data visualizations	Available
Image generation	Available
Exporting a conversation (Markdown, PDF)	Available

Your content and data

Feature	Availability	Notes
Dashboards	Available	Browse dashboards natively and open them in a viewer.
Activity feed	Available	Includes your calendar and to-dos.
Library (saved artifacts)	Available	Includes documents and visuals generated in Quick.
Knowledge graph (view and manage)	Not available	The knowledge graph is on by default, but viewing and managing it is not available on mobile.
Memory (view and manage)	Not available	Quick uses the same memories across desktop and mobile, but managing them on mobile is not yet available.

Connectors, agents, and skills

Feature	Availability	Notes
Agents	Available	Choose an agent and use agent templates.
Connectors	Limited	You can connect a service and re-authenticate an existing connection. Set up, manage, and remove connectors on desktop.
Skills	Limited	Skills run in your conversations. Manage skills on desktop.
Mission Control	Available	
Apps (Apps in Amazon Quick)	Not available	

Settings

Feature	Availability	Notes
Account, language, region, and appearance	Available	
Biometric app lock	Available	
Default agent and response preferences	Available	

Mobile-specific behaviors

Some capabilities work differently on mobile than on desktop.

- Dashboards open in a viewer: on mobile, you browse your dashboards in a native list and open a dashboard in an embedded viewer with mobile controls.
- On-device file access: Quick can read and write files on your device when you grant access. You approve each request, and you control which folders Quick can use.
- Approvals as cards: when an agent needs your permission to take an action, the request appears as a card in the conversation that you can allow or deny.
- Biometric app lock: when biometric lock is on, Amazon Quick requires Face ID, Touch ID, or a fingerprint each time you open the app.

What is not available on mobile

The following capabilities are available on desktop but not on mobile. Use the desktop application for these tasks.

- Apps (Apps in Amazon Quick): build and run apps on desktop.
- Memory management: view and manage your memories on desktop.
- Knowledge graph: view and manage your knowledge graph on desktop.
- Incognito mode: start a private chat on desktop.

- Connector management: set up, manage, and remove connectors on desktop. On mobile, you can connect a service and re-authenticate an existing connection.
- Skill management: manage your skills on desktop. On mobile, skills run in your conversations.

Differences between iOS and Android

The mobile experience is consistent across iOS and Android, with a few differences.

- Mission Control: available on iOS. Not available on Android.
- Branching conversations: available on Android. Not available on iOS.
- Attachments: on iOS, you can attach up to three files per message.
- Creating a connector: on iOS, creating a new connector requires iOS 17.4 or later. You can re-authenticate an existing connector on any supported version.

Document history for the Amazon Quick User Guide

This page describes changes to the *Quick User Guide*. For notifications about these documentation changes, subscribe to the RSS feed using the link near the top of this page.

To find out what's new in Quick, see [What's New with Analytics](#). To find out about changes to the Quick service, sign up for the [Quick newsletter](#).

New Amazon Quick releases appear in different AWS Regions over a period of time, beginning with the first Region on the initial release date. Documentation is released in coordination with this process. If you have questions, contact [AWS Support](#) or your technical account manager.

 **Note**

The following table describes important changes in each *Amazon Quick User Guide* release since March 4, 2019.

Change	Description	Date
Amazon Quick on desktop and mobile refresh	Refreshed the Amazon Quick on desktop chapter for the current interface and navigation model (Chat, Feed, Library, Customize, Apps, Spaces, More); added activity feed, Library, Apps in Amazon Quick, Spaces, My context, and Memories topics; reorganized agents, skills, connectors, and knowledge under the Customize hub; and added a new Amazon Quick on mobile chapter. The artifact library was renamed from "My stuff" to Library.	September 5, 2026

[AWS Agent Registry integration](#)

Added a new integration guide that describes how to connect an AWS Agent Registry to Amazon Quick and surface its remote MCP servers as connectors. For more information, see [AWS Agent Registry integration](#).

August 31, 2026

[Pricing and billing detail for standalone admin pages](#)

Expanded the standalone subscription admin pages with pricing and billing detail, including how collaborator billing works, annual subscription terms, billing cycles and proration, upgrade and downgrade billing, and how additional usage is billed. For more information, see [Upgrading your Amazon Quick plan](#), [Plan and billing for Amazon Quick Free, Plus, and Max accounts](#), and [Understanding usage on the Amazon Quick Plus and Max plans](#).

August 28, 2026

[DLP vended logs](#)

Added documentation for the DLP_LOGS vended log type, which captures data loss prevention (DLP) enforcement decisions and configuration changes. For more information, see [Monitoring Amazon Quick using CloudWatch Logs](#).

August 25, 2026

[Okta desktop enterprise sign-in uses the org authorization server](#)

Updated the Okta enterprise sign-in guide for Amazon Quick on desktop to use the Okta org authorization server, which every Okta org includes at no additional cost. This removes the dependency on the API Access Management add-on that a custom authorization server (for example, `/oauth2/default`) requires. Updated the OIDC endpoints, removed the scope configuration step, and added Okta reference links. For more information, see [Set up enterprise sign-in with Okta](#).

August 10, 2026

[New Limit profiles page](#)

Added a new page for Limits Management limit profiles, which let administrators cap per-user consumption of index storage and agent hours. For more information, see [Limit profiles](#).

August 10, 2026

Data loss prevention with Microsoft Purview	Added a new Governance section with a Data loss prevention page that describes how to connect Microsoft Purview and enforce data loss prevention policies on sensitive files in Amazon Quick. For more information, see Data loss prevention .	August 10, 2026
Added approval workflows documentation	Added documentation for approval workflows, which give administrators governance control over sharing actions for knowledge bases, spaces, and custom chat agents. For more information, see Approval workflows .	August 10, 2026
New custom permissions for browser extension and Office add-ins	Added five new custom permission features under the Extensions parent capability: Browser Extension, Excel Extension, Outlook Extension, PowerPoint Extension, and Word Extension. For more information, see Custom permissions .	July 6, 2026
Change email and social sign-up URL	Update links for Free/Plus/Max accounts to aws.com/quick .	June 24, 2026

[Public applications for Apps in Quick](#)

Free and Plus accounts can now create public-facing applications that are accessible to anyone with the URL. Public applications can be configured with access to in-app storage and AI inference . For more information, see [What you can build](#).

June 12, 2026

[Desktop documentation restructured with Capabilities section](#)

Reorganized the Amazon Quick on desktop documentation. Added a new Capabilities section with dedicated pages for Connections, Skills, Agents, Plugins, and Tools. Moved System tools and MCP server configuration under Capabilities. Added Agents documentation for creating, configuring, and managing agents. Added official catalog information to the Skills page.

June 12, 2026

[OneDrive knowledge base admin-managed setup](#)

Added admin-managed connection support for the Microsoft OneDrive knowledge base connector , including service credentials setup with AWS KMS signing, certificate-based authentication, and tenant-wide crawling of all users' OneDrive content. Admin-managed OneDrive knowledge bases always enforce document-level access control (ACL) with real-time verification at query time. For more information, see [Admin-managed setup](#).

June 1, 2026

[Web search processed in-Region for Sydney and Tokyo](#)

Updated cross-Region web search documentation to reflect that web search queries originating in Asia Pacific (Sydney) and Asia Pacific (Tokyo) are now processed in the originating Region and are no longer routed to US East (N. Virginia) . For more information, see [Cross-Region calls for web search](#).

May 29, 2026

[VPC support for MCP integration](#)

You can now connect Amazon Quick MCP integrations to private MCP servers that are reachable from a virtual private cloud (VPC) by selecting a Amazon Quick VPC connection during connector setup. The OAuth endpoints used by the MCP server must remain reachable over the public internet. Added troubleshooting guidance for connector creation, VPC connection issues, and Microsoft Entra ID errors. For more information, see [Model Context Protocol \(MCP\) integration](#).

May 29, 2026

[Updated HubSpot integration documentation](#)

Restructured the HubSpot integration topic to follow the built-in connector format. The topic now covers prerequisites, account verification, account configuration, Amazon Quick setup, and troubleshooting. The HubSpot action connector reads CRM data, including contacts, companies, deals, and tickets, and authenticates with a custom OAuth app. For more information, see [HubSpot integration](#).

May 28, 2026

[Clarified IAM permissions for DeleteAccountSubscription](#)

Clarified that the DeleteAccountSubscription Amazon Quick API requires both quicksight:Unsubscribe and quicksight:DeleteAccountSubscription IAM permissions. An explicit deny on either permission results in an AccessDeniedException. Updated the account termination permissions list and IAM policy examples accordingly. For more information, see [Closing your Amazon Quick account](#).

May 27, 2026

[Preview extensions for Microsoft 365](#)

Added preview extension documentation for Excel, Word, and PowerPoint with simplified deployment through the Microsoft 365 admin center. For more information, see [Connect Amazon Quick to your applications](#).

April 28, 2026

New connectors	Added connector documentation for Airtable, Dropbox, Google Analytics, Google Calendar, Google Docs, Google Drive, Google Meet, Google Sheets, Google Slides, Gmail, QuickBooks, and Zoom. Updated existing Slack and Outlook connectors. For more information, see Connectors .	April 28, 2026
Free and Plus account documentation	Added documentation for Amazon Quick Free and Plus accounts created at quick.aws.com, including signup, plans, user management, billing, and account administration. For more information, see Signing up at quick.aws.com and Administering Amazon Quick (Free and Plus) .	April 28, 2026
Document and visual creation	Added documentation for creating documents and visuals including Word, Excel, PowerPoint, PDF, and infographics directly from chat conversations. For more information, see Document and visual creation .	April 28, 2026

[Build web applications with apps in Amazon Quick](#)

Added documentation for apps in Amazon Quick, a feature that lets you build, publish, and share interactive web applications using natural language. For more information, see [Build web applications with apps in Amazon Quick](#).

April 28, 2026

[Amazon Quick on desktop](#)

Added documentation for the Amazon Quick desktop application for macOS and Windows, including local file access, activity feed, scheduled agents, system tools, MCP servers, knowledge graph, and settings. For more information, see [Amazon Quick on desktop](#).

April 28, 2026

[Visier integration](#)

Added documentation for the Visier people analytics integration, which connects to Visier through MCP server connectivity. You can query workforce data, analyze headcount trends, assess retention health, and generate workforce briefings. For more information, see [Visier integration](#).

April 22, 2026

[Spaces file size and storage limit increase](#)

Increased the file upload size limit for Spaces to 500 MB for PDF, Word, and PowerPoint files (50 MB for other file types), and increased the total space storage limit from 1 GB to 10 GB. For more information, see [Working with Spaces](#).

April 21, 2026

[Confluence knowledge base document-level access controls](#)

Added document-level access control (ACL) support for the Atlassian Confluence Cloud knowledge base connector. Includes Atlassian admin credentials setup, real-time permission verification at query time, and end-user sign-in flow. For more information, see [Document-level access controls](#).

April 21, 2026

[SharePoint knowledge base admin-managed setup](#)

Added admin-managed connection support for the SharePoint Online knowledge base connector, including service credentials setup (2LO) with AWS KMS signing, certificate-based authentication, document-level access control (ACL) with real-time verification, Sites.Selected least-privilege permissions, and specific SharePoint path support. For more information, see [Admin-managed setup](#).

April 16, 2026

[Knowledge base and data source sharing](#)

Added support for sharing admin-managed SharePoint and Google Drive knowledge bases with co-owners, and sharing data source connections with other users. For more information, see [Sharing knowledge bases and data sources](#).

April 16, 2026

[Index usage monitoring](#)

Added index usage logs (INDEX_USAGE_LOGS) to Amazon Quick vended logs for tracking storage consumption across knowledge bases and Spaces. Includes CloudWatch Log Insights queries for building dashboards. For more information, see [Monitor index storage usage](#).

April 16, 2026

[Incremental refresh support for datasets in the new data preparation experience](#)

You can now incrementally refresh SQL-based SPICE datasets that were created in the new data preparation experience. For more information, see [Incrementally refreshing a dataset](#).

April 15, 2026

[ServiceNow integration documentation update](#)

Rewrote the ServiceNow integration documentation with detailed setup instructions for both user authentication (OAuth authorization code) and service authentication (client credentials), a complete list of available actions derived from the ServiceNow REST API, and a new limitations and troubleshooting section. For more information, see [ServiceNow integration](#).

March 30, 2026

[Fine-grained ACL support for Google Drive knowledge base](#)

Added admin-managed connection support for the Google Drive knowledge base connector, including Google Workspace configuration with 2-legged OAuth (2LO) and document-level access control list (ACL) crawling. For more information, see [Google Drive knowledge base](#).

March 30, 2026

[Updated Quick Research documentation](#)

Updated Quick Research documentation to reflect current UI changes, including research mode options (Fast and Deep), updated research objective guidance, revised button labels, and added third-party data sources as a research material type. For more information, see [Using Amazon Quick Research](#).

March 26, 2026

[New Relic integration](#)

Added New Relic integration documentation with setup instructions, available actions, and troubleshooting. For more information, see [New Relic integration](#).

March 26, 2026

[Region expansion for Tokyo, Frankfurt, and London](#)

Updated regions documentation to reflect full Amazon Quick support in Tokyo (ap-northeast-1), Frankfurt (eu-central-1), and London (eu-west-2). Consolidated cross-Region inference (CRIS) documentation into the regions page with updated inference boundaries for AU-CRIS, JP-CRIS, EU-CRIS, and US-CRIS geographies, and added web search regional availability details. For more information, see [Supported Regions](#).

March 24, 2026

[New custom permissions for Spaces and Chat Agents](#)

Added three new custom permission features: Create Spaces, Share Spaces, and Share Chat Agents. For more information, see [Custom permissions](#).

March 24, 2026

[New Amazon CloudWatch metrics for Amazon Quick index capacity](#)

Added two new Amazon CloudWatch metrics for monitoring Amazon Quick index billing: QuickIndexCapacityConsumedRawFileSizeInGB and QuickIndexCapacityRawFileSizeLimitInGB. For more information, see [Monitoring data in Amazon QuickSight](#).

March 24, 2026

[Google Drive integration](#)

Refactored the Google Drive knowledge base integration documentation

March 20, 2026

[Embed custom chat agents](#)

Added documentation for embedding custom chat agents on external websites, including the process for allowing hosting domains and generating embed code. For more information, see [Create, customize, and deploy AI-powered chat agents](#).

March 19, 2026

[Custom permissions update](#)

Updated custom permissions documentation with new parent capabilities (Scenarios, Stories, Topics), new features (Edit Visual with AI, Build Calculation with AI, Create Executive Summary, Knowledge Base create/update and sharing), and an Action Connectors section documenting per-connector permissions. For more information, see [Custom permissions](#).

March 19, 2026

[Quick Automate actions update](#)

Updated Amazon Quick Automate documentation for the Data tables, Excel, Browser, Custom Agent, REST API, Code, and AWS Integrations action categories.

March 18, 2026

[Atlassian Jira Cloud integration](#)

Added dedicated Atlassian Jira Cloud integration documentation with validated OAuth scopes and a complete actions table. For more information, see [Atlassian Jira Cloud integration](#).

March 18, 2026

[Atlassian Confluence Cloud integration](#)

Refactored the Atlassian Confluence Cloud integration documentation with shared Atlassian OAuth modules for a consistent setup experience.

March 17, 2026

[Namespace isolation and administrative console clarification](#)

Clarified the difference in permission models between the Amazon Quick end-user console and the administrative console for namespace isolation. For more information, see [Supporting multitenancy with isolated namespaces](#).

March 16, 2026

[MCP integration documentation](#)

Rewrote the Model Context Protocol (MCP) integration topic with an MCP protocol overview, OAuth discovery flow documentation (RFC 9728), and shared Dynamic Client Registration (DCR) concepts. For more information, see [Model Context Protocol \(MCP\) integration](#).

March 16, 2026

[Web Crawler VPC support](#)

Added VPC support, sitemap discovery, and JavaScript navigation troubleshooting to the web crawler integration documentation. For more information, see [Web Crawler integration](#).

March 13, 2026

[Managing shared folders custom permission](#)

Added documentation for the managing shared folders custom permission. For more information, see [Custom permissions](#).

March 12, 2026

Chat user preferences	Added a user preferences section to the chat documentation for managing display and notification preferences in Amazon Quick. For more information, see Ask questions, explore data, and get insights with chat .	March 11, 2026
Microsoft 365 integration guides	Rewrote the Microsoft Outlook, SharePoint, and Teams integration guides with a consistent structure that includes full Microsoft Entra setup instructions, permission tables with descriptions, and correct OAuth scopes.	March 10, 2026
Conversation retention period update	Updated the chat conversation retention period from 30 days to 90 days. For more information, see Ask questions, explore data, and get insights with chat .	March 5, 2026
Amazon S3 ACL documentation	Added document-level access control list (ACL) documentation for the Amazon S3 integration and a best practices guide for managing ACLs in knowledge bases. For more information, see Document-level ACLs .	March 2, 2026

Asia Pacific (Malaysia) Region support	Added the Asia Pacific (Malaysia) ap-southeast-5 Region to the list of supported Amazon Quick Sight Regions, including the API endpoint and IP address range. For more information, see Supported Regions .	March 1, 2026
Quick Flows vs Quick Automate comparison	Added a comparison section to the automations landing page to help users choose between Amazon Quick Flows and Amazon Quick Automate based on their workflow needs. For more information, see Choosing between Quick Flows and Quick Automate .	February 27, 2026
TLS 1.2 MySQL version requirement	Documented the requirement for MySQL version 5.7.28 or later to support TLS 1.2 connections for MySQL data sources.	February 26, 2026
ApproveFlowShareRequests custom permission	Added the ApproveFlowShareRequests custom permission, which controls whether users can approve requests to share flows. For more information, see Custom permissions .	February 26, 2026
Space extracted text limits	Updated the extracted text limits for CSV and Excel files in spaces.	February 23, 2026

Knowledge base file size limit increase	Increased the file size limit for standard text documents in knowledge bases from 50 MB to 500 MB.	February 16, 2026
Automations usage metrics	Added a section for tracking Quick Flows, Quick Automate, and Amazon Quick Research usage hours in the usage metrics documentation.	February 6, 2026
SPICE capacity updates	Updated SPICE capacity and dataset size documentation to reflect increased limits.	February 2, 2026
Embedding and row-level security guidance	Added security guidance for embedded analytics to warn against AllowedEmbeddingDomains bypass vulnerabilities. Added guidance for using row-level security (RLS) in multi-tenant architectures.	January 27, 2026
CSV and Excel file size limits	Updated the file size limits for CSV and Excel uploads to 10 MB.	January 27, 2026
Extension access with Okta and IAM Identity Center	Updated extension setup instructions for Microsoft Word, Slack, and Teams to support Okta as an identity provider with IAM Identity Center. For more information, see Extension access .	January 26, 2026

AgentHours vended logs	Updated Quick vended logs documentation to include the AgentHours log type for tracking agent usage.	January 6, 2026
New integrations and PagerDuty advance	Added individual integration pages for 10 new MCP-based services and PagerDuty advance integration. Updated the supported integrations table and connector APIs. For more information, see Model Context Protocol (MCP) integration .	January 5, 2026
User-driven license upgrades	Added documentation for user-driven license upgrades, which allow end users to request license tier changes. For more information, see User-driven license upgrades .	December 29, 2025
Flows in browser extension	Added documentation for using Amazon Quick Flows in the browser extension. For more information, see Using Amazon Quick Flows in browser extension .	December 15, 2025

[Add IAM Identity Center support for extensions](#)

Added documentation for adding Amazon Quick extensions to Microsoft Word, Microsoft Office, Microsoft Teams, and Slack with accounts that use IAM Identity Center. For more information, see [Extension access](#).

November 19, 2025

[New data preparation experience](#)

Added documentation for the new visual data preparation experience in Amazon Quick Sight. For more information, see [Preparing data with the new data preparation experience](#).

October 20, 2025

[New Amazon Quick capabilities](#)

October 9, 2025

Amazon Quick provides an integrated suite of AI-powered analytics, automation, and research tools that transcend traditional business intelligence.

Chat and AI capabilities:

- Amazon Quick chat - Natural language interface for data exploration with contextual awareness, file uploads, web search, and artifact generation
- Quick chat agents - Customizable AI assistants with domain-specific expertise, configurable personas, and organizational context

Workflow automation:

- Amazon Quick Flows - Intelligent workflows combining generative AI with automated actions through no-code interface, supporting conversational and structured runtime modes
- Amazon Quick Automate - AI-powered automation platform that creates sophisticated workflows

from natural language
descriptions or process
documents

Research and analysis:

- Amazon Quick Research - Comprehensive research capability that analyzes multiple data sources and generates detailed reports with citations and source tracing for market research, competitive analysis, and business intelligence

Data and knowledge management:

- Amazon Quick Index - Data discovery and cataloging capabilities
- Spaces - Collaborative workspaces for organizing and sharing data, dashboards, topics, and files with team-specific knowledge bases

Integration and connectivity:

- Integrations - Connect to external applications for unstructured data access, including Google Drive,

Amazon S3, OneDrive, and MCP servers

- Connectors - Perform actions in third-party services like Microsoft Outlook, PagerDuty, and Slack with 15+ supported integrations
- Knowledge bases - Create AI-accessible content from third-party applications and data sources
- Extensions - Browser, Slack, Microsoft Teams, Outlook, and Word extensions for AI assistance within existing tools
- Bring Your Own Index (BYOI) - Connect existing Amazon Q Business indexes as data sources

[Amazon QuickSight rebrands to Amazon Quick](#)

Amazon QuickSight has been rebranded to Amazon Quick, expanding from a standalone business intelligence service to a comprehensive analytics and AI platform. QuickSight continues as Amazon QuickSight, a core component within the Quick ecosystem.

October 9, 2025

Amazon QuickSight retains existing functionality including interactive data visualization, SPICE in-memory analytics, embedded analytics, and dashboard sharing. All existing QuickSight APIs, SDKs, and integrations continue to work without changes.

[Font customization](#)

Authors can now customize fonts for data labels and axes in supported charts, in addition to the previously supported font customizations for visual titles, subtitles, and legends, as well as table and pivot table headers.

October 6, 2025

[Google Sheets data source](#)

You can now create datasets using a Google Sheets data source. Admins will need to perform a one-time setup to enable Google Sheets as a data source.

September 27, 2025

Q&A null support	Q&A now supports comprehensive support for null value handling, enabling users to create more sophisticated analyses and answer complex business questions.	September 18, 2025
Interactive Highchart features	Highchart visualizations now support custom actions, highlighting, and custom field color consistencies, allowing you to create interactive and visually cohesive charts that integrate seamlessly with other visuals.	September 2, 2025
Calculated fields limitation increase	The number of limits for calculated fields in an analysis has increased from 200 to 500. For more information, see Item limits for analyses .	August 15, 2025
Apache Impala data source	You can now create datasets using an Apache Impala data source.	August 1, 2025
Improve Q&A accuracy with custom instructions	Authors can now curate responses to questions by adding domain-specific knowledge that cannot be captured through a topic's metadata settings. This ensures more accurate, relevant, and tailored answers that are better suited for business needs.	July 18, 2025

New topic: Trusted identity propagation support with Athena data sources	Added information about using trusted identity propagation with QuickSight and Athena. For more information, see Using trusted identity propagation with Athena .	June 30, 2025
2 billion row SPICE datasets	QuickSight now supports SPICE datasets that contain up to 2 billion (2,000,000,000) rows. For more information, see SPICE quotas for imported data .	June 30, 2025
Topic Permissions	Authors can now manage permissions for dashboards and their linked topics from a single location. Control topic viewer access directly from a dashboard's sharing preferences, eliminating the need to manage permissions in multiple locations.	April 18, 2025
Sheet Highlighting	Added the ability to highlight specific data points across visuals in a sheet. When you select or hover over a data point on a visual, related data across other visuals will stand out, while unrelated data is dimmed.	April 11, 2025

Updated topic: Opting out of Amazon Q in Quick	Added additional information and guidance for opting out of Amazon Q in Quick. For more information, see Opting out of Amazon Q in Quick .	April 8, 2025
Enable Generative BI features in embedded dashboards and consoles	Added documentation for adding executive summaries to embedded dashboards, and adding executive summaries, data stories, authoring, and Q&A to embedded consoles. For more information, see Embedding with the Amazon QuickSight APIs .	March 31, 2025
Customize alert emails	Previously, the custom email template in your QuickSight account only affected dashboard report and pixel perfect report emails. Now, it also customizes alert emails. For more information, see Customizing report and alert emails .	March 31, 2025
Rules Dataset flag	A Rules Dataset flag has been implemented to distinguish permission datasets used for row-level security from regular datasets.	March 28, 2025

Publish a previous version of a dashboard	You can now revert back to a previous version of a dashboard by searching for it under the dashboard's Version History and publishing it. Each dashboard can store up to 1000 versions that are never deleted.	March 28, 2025
Scenarios general availability	Updated the scenarios documentation to support its general availability release. For more information, see Working with scenarios in Amazon QuickSight .	March 25, 2025
New Region	Amazon QuickSight is now available in the Europe (Spain) (eu-south-2) Region. For more information, see Supported AWS Regions for Amazon QuickSight .	March 11, 2025
Dashboard Q&A experience	Amazon QuickSight authors can now to turn on the Q&A experience directly from their dashboards in one click without the need to create a topic in QuickSight. For more information, see Turn on the Dashboard Q&A experience in Amazon QuickSight .	January 31, 2025

Unique key configuration	Amazon QuickSight authors can now configure a unique key column to a QuickSight dataset during data preparation. For more information, see Adding a unique key to an Amazon QuickSight dataset .	December 19, 2024
Amazon QuickSight scenarios	Amazon QuickSight Pro users can now use Amazon Q in Quick to create scenarios to analyze complex business problems with simple natural language. For more information, see Working with scenarios in Amazon QuickSight .	December 4, 2024
Augmenting Amazon QuickSight insights with Amazon Q Business	Amazon QuickSight account admins can now connect their QuickSight account to Amazon Q Business to augment insights with unstructured data sources. For more information, see Augmenting Amazon QuickSight insights with Amazon Q Business .	December 3, 2024
Prompted reports	Amazon QuickSight authors can now create prompts on pixel-perfect reports to allow dashboard users to filter data in on-demand and scheduled reports. For more information, see Setting up prompts for pixel perfect reports .	November 22, 2024

New visual type: layer map	Amazon QuickSight now supports layer map visuals. For more information, see Using layer maps .	November 22, 2024
New visual type: image component	You can now add images to QuickSight dashboards and analyses with image component visuals in Amazon QuickSight. For more information, see Using image components .	November 22, 2024
New visual type: Highcharts visual	Amazon QuickSight authors can now configure Highcharts visuals in QuickSight. For more information, see Using Highcharts .	November 22, 2024
Import visual public preview	(Public preview) Amazon QuickSight authors can now import QuickSight visuals from one dashboard or analysis to a new analysis that has access permissions. For more information, see Importing Amazon QuickSight visuals to an analysis .	November 22, 2024
Font customization	Amazon QuickSight dashboard viewers can now schedule their own dashboard reports for themselves. For more information, see Creating a reader generated report in Amazon QuickSight .	November 22, 2024

Font customization	Amazon QuickSight authors can now customize the font, font size, and font color of visual titles, subtitles, legends, headers, and cells. For more information, see Titles and subtitles on visual types in QuickSight .	November 22, 2024
Cross-Region inference with Amazon Q in Quick	Amazon Q in Quick now supports cross-Region inferencing. For more information, see Cross-Region inference with Amazon Q in Quick .	November 21, 2024
Create a custom brand in Amazon QuickSight	Amazon QuickSight account admins can now customize their application's branding and visual theme to align with their organization's guidelines. For more information, see Amazon QuickSight brand customization .	November 14, 2024
Connect to Starburst with OAuth client credentials	You can now use OAuth client credentials to connect your Amazon QuickSight account with Starburst through the QuickSight APIs. For more information, see Creating an Amazon QuickSight data source connection to Starburst with OAuth client credentials .	November 8, 2024

[Connect to Snowflake with OAuth client credentials](#)

You can now use OAuth client credentials to connect your Amazon QuickSight account with Snowflake through the QuickSight APIs. For more information, see [Creating an Amazon QuickSight data source connection to Snowflake with OAuth client credentials](#).

November 8, 2024

[Subfolders for restricted shared folders](#)

Assets that are located in a QuickSight restricted folder can now be moved within the restricted folder tree into one or more subfolders. For more information, see [Organizing assets into folders for Amazon QuickSight](#).

October 15, 2024

[Show an apply button](#)

QuickSight authors can now add an **Apply** button to a Calendar, Relative date picker, or Multiselect dropdown control that delays visual reload until the user chooses **Apply**. For more information, see [Adding filter controls to analysis sheets](#).

September 27, 2024

[Data story personalization](#)

Amazon Q in Quick now leverages user location and job-related information from your IAM Identity Center instance to generate personalized data stories that are more relevant to authors and readers. For more information, see [Personalize data stories in Amazon QuickSight](#).

September 27, 2024

[Shared view for embedded dashboards](#)

QuickSight developers can now use the Amazon QuickSight Embedding SDK (version 2.8.0 and higher) to allow readers of embedded dashboards to receive and distribute shareable links to their view of an embedded dashboard. For more information, see [Share an embedded view](#).

August 23, 2024

[5 new regions](#)

Amazon Q in Quick is now available in 5 new opt-in Regions: Asia Pacific (Mumbai), Canada (Central), Europe (Ireland), Europe (London), and South America (São Paulo). For more information, see [Supported AWS Regions for Amazon Q in Quick](#).

August 9, 2024

Nested filters	You can now apply nested filters to QuickSight analyses. For more information, see Adding nested filters .	August 1, 2024
Joined datasets in SPICE	The size limit for secondary tables of joins for datasets that are stored in SPICE have been updated. For more information, see Joining data .	July 10, 2024
Advanced options for paginated reports	You can now create repeating sections for pixel perfect reports in Amazon QuickSight. For more information, see Create repeating sections .	June 27, 2024
4 New Regions	QuickSight is now available in 4 new opt-in Regions: Europe (Milan), Europe (Zurich), Africa (Cape Town), and Asia Pacific (Jakarta). For more information, see AWS Regions, websites, IP address ranges, and endpoints .	May 28, 2024
New sorting options for tables in Amazon QuickSight	QuickSight Authors can now configure single column, multi column, and off-visual sorts for tables in Amazon QuickSight. For more information, see Sorting tables .	May 20, 2024

[Key management in
QuickSight: Updated IAM
policy conditions](#)

The required IAM policy permissions that grant access to AWS KMS keys in the QuickSight console and the QuickSight APIs have been updated. For more information, see [IAM identity-based policies for Amazon QuickSight: using the admin key management console](#).

May 16, 2024

[???](#)

QuickSight Admins can now update and manage CMKs that are registered to a QuickSight account with the QuickSight APIs. For more information about creating and managing CMKs with the QuickSight APIs, For more information, see [Key management operations](#).

May 16, 2024

[Hide SPICE capacity labels](#)

QuickSight Admins can now choose to hide the account-wide SPICE capacity usage and remaining size labels from QuickSight authors. For more information, see [Hiding SPICE capacity labels](#).

May 2, 2024

[Embed the Generative Q&A experience](#)

You can now set up an embedded Q&A that uses enhanced NLQ capabilities powered by LLMs. The Q&A is the recommended replacement for the embedded Q Search Bar and provides an updated BI experience for users. For more information, see [Embed the Amazon Q in Quick Generative Q&A experience](#).

April 30, 2024

[Amazon Q in Quick: New Pro roles](#)

Amazon Q integrates with Amazon QuickSight to give QuickSight users access to a suite of new generative BI capabilities. With Amazon Q in Quick, you can utilize the generative BI authoring experience, create executive summaries of your data, ask and answer questions of data, and generate data stories. To get started with Amazon Q in QuickSight Generative BI capabilities, upgrade your account's users to Admin Pro, Author Pro, or Reader Pro roles. For more information, see [Get started with Generative BI](#).

April 30, 2024

[VPC endpoint restrictions](#)

You can now create VPC endpoint restrictions in Amazon QuickSight to limit access to your QuickSight account. For more information, see [Turning on IP and VPC endpoint restrictions in Amazon QuickSight](#).

April 9, 2024

[Account instances of IAM Identity Center for account management](#)

Amazon QuickSight now supports account instances of IAM Identity Center for identity management. For more information, see [Managing user access inside Amazon QuickSight](#).

April 9, 2024

[Managed policy update](#)

Amazon QuickSight added a new managed policy to simplify Asset Bundle export operations. For more information, see [AWS managed policy: AWSQuickSightAssetBundleExportPolicy](#).

March 27, 2024

[Managed policy update](#)

Amazon QuickSight added a new managed policy to simplify Asset Bundle import operations. For more information, see [AWS managed policy: AWSQuickSightAssetBundleImportPolicy](#).

March 27, 2024

New Q pane	Amazon QuickSight users that are opted in to the Generative BI preview can now use the Q pane to access all relevant Generative BI features. For more information, see Using Generative BI with Amazon Q in Quick .	March 25, 2024
New Region: China (north)	Amazon QuickSight is now available in China (Beijing) (cn-north-1). For more information, see AWS Regions, websites, IP address ranges, and endpoints .	January 25, 2024
Amazon Q Business in Amazon QuickSight	Amazon Q Business integrates with Amazon QuickSight to give QuickSight users access to a suite of new generative BI capabilities. With Amazon Q in Quick, you can utilize the generative BI authoring experience, create executive summaries of your data, ask and answer questions of data, and generate data stories. For more information, see Using generative BI with Amazon Q in Quick .	November 28, 2023

[Trusted identity propagation in Amazon QuickSight](#)

QuickSight authors can now create data sources with trusted identity propagation, which allows database administrators to centrally manage data security in Amazon Redshift and automatically apply all data security rules to data consumers in QuickSight. For more information, see [Enabling trusted identity propagation with Amazon Redshift](#).

November 26, 2023

[Runtime filtering and theming for embedded dashboards and visuals](#)

QuickSight authors can now customize filters and themes for embedded dashboards and visuals. For more information, see [Customizing embedded assets](#).

November 17, 2023

[Custom time zones and week start days in Amazon QuickSight](#)

QuickSight authors can now restrict the functionality that people can access in Amazon QuickSight. For more information, see [Customizing access to the Amazon QuickSight console](#).

November 17, 2023

[Custom time zones and week start days in Amazon QuickSight](#)

QuickSight authors can now turn on SPICE auto capacity purchasing to allow Amazon QuickSight to automatically manage your QuickSight account's SPICE capacity. For more information, see [Turning SPICE auto capacity purchasing on](#).

November 17, 2023

[Amazon EventBridge integration](#)

You can now integrate your Amazon QuickSight account with Amazon EventBridge. For more information, see [Amazon QuickSight EventBridge integration](#).

November 17, 2023

[New look for the QuickSight analysis workspace](#)

We've redesigned the Amazon QuickSight analysis workspace. You might encounter screenshots or procedural text that doesn't reflect the new look in the QuickSight console. We're in the process of updating screenshots and procedural text. For more information on QuickSight's new look, see [Introducing new analysis experience on Amazon QuickSight](#).

November 16, 2023

[Custom time zones and week start days in Amazon QuickSight](#)

QuickSight authors can now set custom time zones and week start days for their analyses. For more information, see [Customize date and time values of an analysis](#).

November 14, 2023

[Float decimal type](#)

Decimal values in calculated fields of datasets that are stored in SPICE can now be assigned the Fixed or Float decimal data type. For more information, see [Supported data types and values](#).

November 8, 2023

[Custom total values in Amazon QuickSight](#)

QuickSight authors can now define the total and subtotal aggregations for their table or pivot table visuals from the field wells. For more information, see [Custom total values](#).

October 25, 2023

[Restricted and unrestricted shared folders in Amazon QuickSight](#)

QuickSight authors can now create restricted and unrestricted shared folders with the QuickSight CLI to help organize and manage the permissions of important QuickSight assets. For more information, see [Organizing assets into folders for Amazon QuickSight](#).

October 24, 2023

New data source	QuickSight now supports connecting to a Trino data source. For more information, see Using Trino with Amazon QuickSight .	October 20, 2023
New data source	QuickSight now supports connecting to a Starburst data source. For more information, see Using Starburst with Amazon QuickSight .	October 20, 2023
SageMaker AI Canvas integration	QuickSight authors can now export data into SageMaker AI Canvas to build ML models that can be sent back to QuickSight and applied to analyses and dashboards. For more information, see Build predictive models with SageMaker AI Canvas .	October 6, 2023
Managed policy update	The AWSQuickSightSageMakerPolicy is updated to reflect the Amazon QuickSight–SageMaker AI integration. For more information, see AWS managed policy: AWSQuickSightSageMakerPolicy .	October 6, 2023

Amazon Redshift run-as role	QuickSight authors can now connect Amazon Redshift data with run as role to enhance data security with fine grained access policies. For more information, see Run queries as an IAM role in Amazon QuickSight .	October 6, 2023
Generative BI public preview	The public preview for Generative BI is now available to users who have a QuickSight Q add-on subscription. For more information, see Working with Generative BI in QuickSight Q .	September 28, 2023
Freeze table columns in Amazon QuickSight	You can now freeze single or groups of table columns in Amazon QuickSight. For more information, see Freeze columns to table visuals .	September 26, 2023
Enhanced KPI options in Amazon QuickSight	New layouts and formatting options for KPIs are now available in Amazon QuickSight. For more information, see Using KPIs .	September 15, 2023
Scaled folders in Amazon QuickSight	You can create scaled folders that can be shared with thousands of namespaces at once with the Amazon QuickSight APIs. For more information, see Creating scaled folders with the AWS CLI .	August 30, 2023

[Excel format reports in Amazon QuickSight](#)

You can now create and schedule Excel snapshot reports of a dashboard in Amazon QuickSight. For more information, see [Scheduling and sending reports by email](#).

August 24, 2023

[Add contextual row subtitles to pivot tables in Amazon QuickSight](#)

You can now add contextual row subtitles to pivot tables in Amazon QuickSight. For more information, see [Totals and subtotals](#).

August 16, 2023

[New layouts for pivot tables in Amazon QuickSight](#)

Amazon QuickSight now supports two types of layout for pivot tables: tabular and hierarchy. For more information, see [Choosing a layout](#).

August 11, 2023

[Default column width Amazon QuickSight](#)

You can now set the default column width of your pivot table visuals in Amazon QuickSight. For more information, see [Row and column size in tables and pivot tables in Amazon QuickSight](#).

August 11, 2023

[Amazon QuickSight integration with IAM Identity Center](#)

You can now integrate your Amazon QuickSight Enterprise edition account with IAM Identity Center. For more information, see [Configure your Amazon QuickSight account with IAM Identity Center](#).

August 11, 2023

[Embedded callback actions at runtime](#)

You can now use embedded data point callback actions to build tighter integrations between your SaaS application and your Amazon QuickSight embedded dashboards and visuals. For more information, see [Add embedded callback actions to at runtime in Amazon QuickSight](#).

August 9, 2023

[Working with field level coloring in Amazon QuickSight](#)

You can now assign specific colors to specific field values across all visuals in a QuickSight analysis or dashboard. For more information, see [Working with field level coloring in Amazon QuickSight](#).

July 13, 2023

[Small multiples axis options in Amazon QuickSight](#)

You can now configure the x and y axes for each individual panel of a small multiples visual in Amazon QuickSight. For more information, see [Small multiples axis options](#).

July 13, 2023

[Use the analysis menu in Amazon QuickSight](#)

You can now use menu options to efficiently perform tasks without needing to manually navigate through your analysis in Amazon QuickSight. For more information, see [The analysis menu](#).

July 7, 2023

[Advanced report scheduling options in Amazon QuickSight](#)

You can now access more advanced PDF report scheduling options in Amazon QuickSight. For more information, see [Configuring email report settings for a dashboard in QuickSight Enterprise edition](#).

June 30, 2023

[Format visual language data based on language settings in Amazon QuickSight](#)

You can now choose how your numeric data values appear in visuals so that they align with the regional language that you have chosen in Amazon QuickSight. For more information, see [Formatting visual numeric data based on language settings in Amazon QuickSight](#).

May 26, 2023

[Create geospatial heatmaps in Amazon QuickSight](#)

You can now create geospatial heatmaps in Amazon QuickSight. For more information, see [Geospatial heatmaps in Amazon QuickSight](#).

May 26, 2023

[Work with Q topics using the QuickSight CLI in Amazon QuickSight](#)

You can now work with QuickSight Q topics using the Amazon QuickSight command line interface (CLI). For more information, see [Work with QuickSight Q topics using the Amazon QuickSight CLI](#).

May 4, 2023

[Use dataset parameters in Amazon QuickSight](#)

You can now use dataset parameters in direct query to dynamically customize their datasets and apply reusable logic to your datasets. For more information, see [Using dataset parameters in Amazon QuickSight](#).

May 4, 2023

[Enhanced options for scatter plots in Amazon QuickSight](#)

You can now plot unaggregated values on scatter plots in Amazon QuickSight. For more information, see [Using scatter plots](#).

May 4, 2023

[Create and update VPC connections Amazon QuickSight](#)

You can now create and update VPC connections in Amazon QuickSight. For more information, see [Using dataset parameters in Amazon QuickSight](#).

May 4, 2023

[Create tag-based RLS rules using the OR condition in Amazon QuickSight](#)

You can now add the OR condition to your tag-based rules to further customize the way data is presented to your QuickSight account users. For more information, see [Using row-level security \(RLS\) with tag-based rules to restrict access to a dataset when embedding dashboards for anonymous users](#).

April 7, 2023

[Threshold alerts on Table and Pivot table visuals in Amazon QuickSight](#)

You can now create threshold alerts for Table and Pivot table visuals in Amazon QuickSight. For more information, see [Working with threshold alerts in Amazon QuickSight](#).

March 17, 2023

[Hide collapsed columns in Amazon QuickSight table visuals](#)

You can now customize the way Amazon QuickSight readers view pivot tables to make them easier to read and understand at a glance. For more information, see [Pivot table layout options](#).

March 9, 2023

[Enhanced embedding capabilities in Amazon QuickSight](#)

You can now use QuickSight's Embedding SDK (v2.0) to improve developer experience when embedding QuickSight in your application or website. For more information, see [Embedding with the QuickSight APIs](#).

March 9, 2023

[Run queries with Amazon S3 data sources in Amazon QuickSight](#)

You can now enhance data security by using fine-grained access policies rather than broader permissions for data sources connected to Amazon S3 in Amazon QuickSight. For more information, see [Run queries with Amazon S3 data sources](#).

February 21, 2023

[Radar charts in Amazon QuickSight](#)

You can now create radar charts in Amazon QuickSight. For more information, see [Using radar charts in Amazon QuickSight](#).

January 30, 2023

[Data bars for tables and pivot tables in Amazon QuickSight](#)

You can now use data bars to add visual context to your table visuals in Amazon QuickSight. For more information, see [Adding data bars to tables in Amazon QuickSight](#).

January 24, 2023

[New question types: Boolean, Forecast, and Why](#)

You can now enter boolean, forecast, and why questions into the QuickSight Q search bar. For more information, see [Asking questions with Amazon QuickSight Q](#).

November 29, 2022

[Automated data prep for QuickSight Q](#)

Amazon QuickSight Q now uses AI-enhanced data preparation to automatically create topics that are relevant to your end users. For more information, see [Working with Amazon QuickSight Q topics](#).

November 29, 2022

[Pixel perfect reporting](#)

You can now create, schedule, and share highly formatted multi-page reports in Amazon QuickSight. For more information, see [Working with pixel perfect reports in Amazon QuickSight](#).

November 28, 2022

[Embedding the Q search bar for anonymous \(unregistered\) users](#)

You can now embed the QuickSight Q search bar for anonymous (unregistered) users. For more information, see [Embedding the Amazon QuickSight Q search bar for anonymous \(unregistered\) users](#).

November 19, 2022

[Asset management](#)

You can now manage all of the assets on your Amazon QuickSight account in one unified view. For more information, see [Amazon QuickSight asset management](#).

November 19, 2022

[Text boxes](#)

You can now add static and dynamic textual content using text boxes in Amazon QuickSight. For more information, see [Using text boxes](#).

November 18, 2022

[Small multiples](#)

You can now create small multiples in line, bar, and pie chart visuals. For more information, see [Using small multiples](#).

November 18, 2022

[Set a reserved value in QuickSight datasets](#)

You can now set a reserved value to determine the value of the **Select all** value of a dataset in Amazon QuickSight. For more information, see [Setting up parameters in Amazon QuickSight](#).

November 18, 2022

[Run queries as an IAM role in Amazon QuickSight](#)

You can now enhance data security by using fine-grained access policies rather than broader permissions for data sources connected to Athena. For more information, see [Run queries as an IAM role in Amazon QuickSight](#).

November 18, 2022

[Pin Amazon QuickSight Q answers to your pinboard](#)

You can now pin visuals from Q for easy access to frequently asked questions. For more information, see [Pinning visuals in Amazon QuickSight Q](#).

November 18, 2022

[Marker clustering on geospatial map visuals](#)

You can now use marker clustering to improve readability of co-located points on a map. For more information, see [Marker clustering on geospatial point maps in QuickSight](#).

November 18, 2022

[Line chart customization](#)

You can now add customizations to line chart visuals to emphasize what you want readers to focus on: color, line style, and markers. For more information, see [Line and marker styling on line charts in QuickSight](#).

November 18, 2022

[Connecting Amazon QuickSight accounts to Databricks](#)

You can now use Amazon QuickSight to connect to Databricks on AWS. For more information, see [Using Databricks in QuickSight](#).

November 18, 2022

[Account termination protection](#) [n](#)

You can now use account termination protection to prevent someone from deleting a QuickSight user account accidentally or maliciously. For more information, see [Deleting your Amazon QuickSight subscription and closing your account.](#)

November 18, 2022

[1 billion row SPICE datasets](#)

Amazon QuickSight now supports SPICE datasets that contain up to 1 billion (1,000,000,000) rows or 1 terabyte (TB) of data. For more information, see [SPICE quotas for imported data.](#)

November 18, 2022

[Monitor SPICE consumption metrics in Amazon CloudWatc](#) [h](#)

You can now monitor your QuickSight account's SPICE consumption metrics in Amazon CloudWatch. For more information, see [Aggregate SPICE metrics.](#)

November 8, 2022

[Use AWS KMS to encrypt your SPICE datasets in Amazon QuickSight](#)

You can now encrypt your SPICE datasets using the keys that you have stored in AWS KMS. This provides you with the tools to audit access to data and satisfy regulatory security requirements. For more information, see [Key management.](#)

October 27, 2022

[Add datasets that contain row-level security \(RLS\) to Q topics](#)

QuickSight Q now supports questions for access restricted datasets that use Row level Security (RLS) with user based rules. Readers can now ask questions about Topics that contain restricted access datasets and instantly receive accurate and pertinent answers based on access control rules defined by authors in RLS settings. For more information, see [Adding datasets with row-level security \(RLS\) to a topic](#).

October 10, 2022

[Use AWS Secrets Manager secrets instead of database credentials in Amazon QuickSight](#)

Amazon QuickSight administrators can grant QuickSight read-only access to secrets they create in AWS Secrets Manager. These secrets can be used in place of database credentials when creating and editing data sources using the QuickSight API. For more information, see [Using AWS Secrets Manager secrets instead of database credentials in Amazon QuickSight](#).

October 6, 2022

[Visual Embedding](#)

You can now embed visuals using a 1-click embed code or with the Amazon QuickSight API. For more information, see [Embedding visuals with the Amazon QuickSight API](#).

August 25, 2022

Undo and redo changes in Q answers	You can now undo or redo any changes you make to a Q answer by choosing the undo or redo arrows in the Q search bar. For more information, see step six in Asking questions with Q .	August 22, 2022
Updates to the sign in experience	The sign in experience in QuickSight has been updated. For more information about how to sign in, see Signing in to QuickSight .	July 21, 2022
Bookmarks	You can now bookmark views of a dashboard. For more information, see Bookmarking views of a dashboard .	July 21, 2022
Base Maps	When creating map visuals in QuickSight, you can now change map base. For more information, see Changing base maps .	July 21, 2022
Q updates	You can now try QuickSight Q out for free before getting the Q add-on. For more information, see Trying QuickSight Q .	July 7, 2022
Rolling dates in filters and parameters	You can now set a rolling date on date range filters and relative date parameters in analyses. For more information, see Creating range date filters in analyses and Setting up parameters .	July 5, 2022

[Calculations update](#)

You can now create level-aware calculations in QuickSight. You can specify at what level to group computations using level-aware calculation - aggregation (LAC-A) functions. You can also specify the window or partition to compute a calculation using level-aware calculation - window (LAC-W) functions (previously known as level-aware aggregations). For more information, see [Using level-aware calculations](#).

July 5, 2022

[Allowed domains](#)

Developers can now add a domain at run-time to an allow list with the `AllowedDomains` parameter of a `GenerateEmbedUrlForAnonymousUser` or a `GenerateEmbedUrlForRegisteredUser` API call. For more information, see [Allow listing domains at run-time with the QuickSight API](#).

July 5, 2022

[Account provisioning](#)

You can now sign up for an Amazon QuickSight account with the Amazon QuickSight API. For more information, see [CreateAccountSubscription](#).

July 5, 2022

[Monitoring data with Amazon CloudWatch](#)

You can now monitor metrics for Amazon QuickSight dashboards, visuals, and dataset ingestions on Amazon CloudWatch to provide your readers with a consistent, high-performing, and uninterrupted experience on QuickSight. For more information, see [Monitoring data in Amazon QuickSight](#).

June 22, 2022

[Forecast boundaries](#)

Authors can now set a minimum and maximum for forecasted values. For more information, see [Forecasting and creating what-if scenarios](#).

June 16, 2022

[Table and pivot table enhancements](#)

Authors and Readers can now resize rows and columns in a table or pivot table visual. For more information, see [Resizing rows and columns in tables and pivot tables](#).

June 14, 2022

[Hiding fields in pivot tables](#)

You can now hide fields in pivot table visuals. For more information, see [Showing and hiding pivot table columns](#).

June 14, 2022

[ML insights and forecasting for Q](#)

Readers can now see key data insights and add forecasts to Q answers. For more information, see [Asking questions with QuickSight Q](#).

June 2, 2022

What's New in QuickSight	To find out what's new in QuickSight, see the What's New page. What's New posts provide a brief overview of all QuickSight feature announcements as they are released.	June 1, 2022
Line chart updates	QuickSight now supports 10,000 data points for line charts. For more information, see Using line charts .	May 22, 2022
1-click public embedding	You can now embed dashboards for anyone on the internet in your public applications with a 1-click static embed code. For more information, see Enabling public access to dashboards with a 1-click embed code .	May 18, 2022
Creating and managing groups	You can now create and manage groups in the QuickSight console. For more information, see Creating and managing groups in Amazon QuickSight .	March 21, 2022
Automatic refresh controls for direct query	In QuickSight, filter controls are now refreshed every 24 hours for direct query. For more information, see Refreshing data .	March 8, 2022

Updating files in a dataset	You can now update files in a dataset to get the latest version of those files. For more information, see Updating files in a dataset .	January 27, 2022
Rich text formatting for titles and subtitles	QuickSight now offers rich text formatting for titles and subtitles, and the ability to add hyperlinks and parameters in titles. For more information, see Formatting a visual title and subtitle .	January 27, 2022
Comparative and cumulative date/time calculations	Comparative and cumulative period functions, such as year-over-year and year-to-date, are now supported in QuickSight. For more information, see Add comparative and cumulative date/time calculations in QuickSight .	January 27, 2022
Sharing dashboards	You now have the option to grant everyone in your QuickSight account access to a dashboard. You can also share a link to the dashboard with anyone who has access to it. For more information, see Sharing dashboards .	November 23, 2021

Customizing email report templates	You can now customize how dashboard email reports appear and behave for QuickSight account users. For more information, see Customizing email report templates .	November 23, 2021
1-click enterprise embedding	You can now embed dashboards for registered users in your internal applications with a 1-click static embed code. For more information, see Embedding dashboards for registered users with a 1-Click embed code .	November 23, 2021
New data source	QuickSight now supports connecting to an Exasol data source. For more information, see Supported data sources .	November 22, 2021
Dataset versioning	QuickSight now supports dataset versioning. For more information, see Reverting datasets back to previous published versions .	November 22, 2021
Adding datasets with RLS to Q topics	QuickSight Q topics now support datasets with row-level security (RLS). For more information, see Adding datasets with RLS to a Q topic .	November 19, 2021

Sheet change performance operations	QuickSight now only refreshes visuals when switching sheets if required. For more information, see Refreshing visuals .	November 12, 2021
Incrementally refresh a SQL-based dataset	In QuickSight Enterprise Edition, you can now refresh your SQL-based SPICE datasets incrementally within a look-back window of time. For more information, see Refreshing data .	October 25, 2021
Embedding the Q search bar	You can now embed the Q search bar in your application for registered users of QuickSight. For more information, see Embedding the QuickSight Q search bar for registered users .	October 22, 2021
Email syncing for federated users	In Enterprise Edition, QuickSight administrators can now restrict new users from using personal email addresses when provisioning through their identity provider (IdP) to QuickSight. For more information, see Configuring email syncing for federated users .	October 22, 2021

Adding RLS tags to a dataset	You can now add tag-based rules to a dataset in QuickSight when applying row-level security (RLS). For more information, see Using row-level security (RLS) with tag-based rules to restrict access to a dataset when embedding dashboards for anonymous users .	October 19, 2021
IP restrictions	You can now limit access to your organization's QuickSight account to a predefined list of Internet Protocol (IP) ranges. For more information, see Turning on Internet Protocol (IP) restrictions in Amazon QuickSight .	October 18, 2021
New table and pivot table styling options	You can now customize the look of tables and pivot tables. For more information, see Formatting tables and pivot tables .	October 12, 2021
New SPICE quota for imported data	In QuickSight Enterprise edition, you can now import up to 500 million rows into a SPICE dataset. For more information, see Data source quotas .	October 12, 2021

Links and images in tables	You can now add links to tables. You can also render links as images in tables. For more information, see Field styling .	October 12, 2021
Custom date formats	You can now customize how dates are formatted in your filter and parameter controls. For more information, see Customizing date formats .	October 8, 2021
QuickSight q	You can now ask questions about your data in the Q bar and get answers in the form of visuals. For more information, see Working with QuickSight q .	September 24, 2021
New geospatial region	QuickSight now supports creating geospatial charts for India. For more information about creating geospatial charts, see Using geospatial charts (maps) .	September 22, 2021
Free-form layouts	You can now place visuals anywhere in your dashboard using precise coordinates. For more information, see Customizing dashboards and visuals .	September 22, 2021
Passing IAM roles to QuickSight	You can now pass IAM roles to QuickSight. For more information, see Using QuickSight with IAM .	September 15, 2021

Create datasets from existing datasets	You can now create datasets from existing datasets. For more information, see Creating a dataset using an existing dataset .	September 15, 2021
Updated AWS managed policy	QuickSight updated an existing AWS managed policy. For more information, see Amazon QuickSight updates to AWS managed policies .	September 8, 2021
New AWS managed policy	QuickSight added a new AWS managed policy. For more information, see Amazon QuickSight updates to AWS managed policies .	September 8, 2021
Forecast multiple measures	You can now forecast up to three measures in the same chart. For more information, see Forecasting and creating what-if scenarios .	August 30, 2021
Share your view of a dashboard	You can now share your view of a dashboard. For more information, see Sharing dashboards .	August 20, 2021
Hide the data zoom on a chart	You can now choose to show or hide the data zoom for your charts in Amazon QuickSight. For more information, see Formatting axes and Grid lines .	August 16, 2021

[Row-level security with tags](#)

You can now use row-level security (RLS) with tags to specify which data your (unauthenticated) users can see in an embedded Amazon QuickSight dashboard depending on who they are. For more information, see [Using RLS with tags to restrict access to a dataset when embedding dashboards for anonymous users](#).

July 29, 2021

[New embedding API operations](#)

Amazon QuickSight has the following new API operations for embedding analytics: `GenerateEmbedUrlForAnonymousUser` and `GenerateEmbedUrlForRegisteredUser`. You can still use the `GetDashboardEmbedUrl` and `GetSessionEmbedUrl` API operations to embed dashboards and the QuickSight console, but they don't contain the latest embedding capabilities. For more information about using the new API operations, see [Embedding overview](#). For more information on using the old API operations, see [Embedding analytics using the old API operations](#).

July 29, 2021

Customizing charts	You can now hide the sort icon and the axis or field title in charts in Amazon QuickSight. For more information, see Formatting axes and Grid lines .	July 16, 2021
Threshold alert updates	You can now choose to be notified when there is no data to check your alert rule against in Amazon QuickSight. For more information, see Working with threshold alerts .	July 9, 2021
New chart type: Custom visual content	You can now embed webpages and online videos, forms, and images in your QuickSight dashboards using the custom visual content chart type. For more information, see Using custom visual content in QuickSight .	June 29, 2021
Feature improvements	You can now show totals for stacked bar charts. For more information, see Using bar charts .	June 29, 2021
Threshold alert updates	Dataset owners can now set a threshold alert evaluation schedule for a dataset in Amazon QuickSight. For more information, see Working with threshold alerts .	June 17, 2021

Duplicating analyses	You can now duplicate an analysis in Amazon QuickSight. For more information, see Duplicating an analysis .	June 17, 2021
Threshold alert updates	You can now choose how often you want to receive threshold alerts in Amazon QuickSight. For more information, see Working with threshold alerts .	May 25, 2021
Wildcard search	Amazon QuickSight now supports wildcard search. For more information, see Using wildcard search .	May 1, 2021
Threshold alerts	Amazon QuickSight now supports creating threshold alerts using KPI and Gauge visuals in dashboards. For more information, see Working with threshold alerts .	May 1, 2021
Service control policies	Amazon QuickSight now supports using service control policies (SCPs) to restrict how individuals in your organization can sign up for QuickSight. For more information, see Using service control policies to restrict QuickSight signup options .	April 26, 2021

Parameter improvements	Amazon QuickSight now supports dynamically displaying parameter values in titles and descriptions throughout charts and analyses. For more information, see Using parameters in titles and descriptions .	April 12, 2021
Feature improvements	QuickSight now supports selecting a time granularity for date and time range filters and datetime parameters. For more information, see Adding a date filter and Setting up parameters .	April 1, 2021
Export to PDF	QuickSight now supports PDF exports of the current sheet in dashboards and analyses. In Enterprise edition, you can now attach a PDF to email reports. For more information, see Exporting an analysis or dashboard as a PDF .	April 1, 2021
Seconds support for date fields	QuickSight now supports aggregating date fields at the second level. For more information, see Changing date field granularity .	March 30, 2021

[Pivot table sorting improvements](#)

QuickSight now supports sorting values in pivots tables by fields in the **Rows** and **Columns** field wells or by column headers. For more information, see [Sorting pivot tables](#).

March 30, 2021

[New feature](#)

Amazon QuickSight now supports customizing tooltips in visuals. For more information, see [Customizing tooltips in a visual](#).

March 16, 2021

[Feature improvements](#)

In Enterprise edition, QuickSight now supports computing anomalies for a time and measure field only, and computing the exact combination of fields in the Category field well. For more information, see [Adding an ML insight to detect outliers and key drivers](#).

March 16, 2021

[Line chart improvements](#)

Amazon QuickSight now supports creating line charts with dual axes. For more information, see [Creating a dual-axis line chart](#).

March 1, 2021

[New formatting options for charts](#)

Amazon QuickSight now supports hiding axis lines, axis labels, and grid lines, and also customizing the how axis labels appear in a chart. For more information, see [Formatting axis lines, axis labels, and Grid lines](#).

February 5, 2021

[New chart type: Sankey diagrams](#)

Amazon QuickSight now supports creating Sankey diagrams. For more information, see [Using sankey diagrams](#).

February 5, 2021

[Field mapping improvements](#)

Amazon QuickSight now supports updating field mapping between mismatched and missing fields when replacing a dataset in an analysis. For more information, see [Replacing datasets](#).

February 5, 2021

[Data prep enhancements](#)

Amazon QuickSight now supports organizing fields into folders and subfolders when preparing or editing datasets. For more information see, [Organizing fields into folders](#).

February 5, 2021

New features

In Enterprise edition, QuickSight now supports displaying anomaly ranges and multiple anomalies when exploring anomalies. Line charts now display notifications when QuickSight detects an anomaly, key driver, or forecasting opportunity. For more information, see [Detecting outliers with ML-powered anomaly detection](#).

February 3, 2021

New region: South America (São Paulo)

Amazon QuickSight is now available in South America (São Paulo) (sa-east-1). For more information, see [AWS Regions, websites, IP address ranges, and endpoints](#).

December 23, 2020

New region: Canada (central)

Amazon QuickSight is now available in Canada (Central) (ca-central-1). For more information, see [AWS Regions, websites, IP address ranges, and endpoints](#).

December 23, 2020

Define an action to open a different sheet

QuickSight now supports navigation actions to enable you to open a different sheet with the included parameter values. For more information, see [Using custom actions for filtering and navigating](#).

December 21, 2020

Data source enhancement: Athena federated query	QuickSight now supports connecting to Athena Federated Query. For more information, see Creating a dataset using Amazon Athena data .	December 21, 2020
New list items sheet control type	Amazon QuickSight now supports single and multi-select list control on dashboards. For more information, see Parameter controls .	December 18, 2020
New formatting options for tables	In QuickSight, you can now reorder columns in table charts. For more information, see Using tables as visuals .	December 17, 2020
Row-level security (RLS) now supports GRANT option only	To streamline and simplify RLS setup, QuickSight no longer supports using RLS in deny-access mode. To create a new RLS configurations, use the explicit grant-access model. Current RLS datasets and deny-access configurations should continue to work as expected. For more information, see Using row-level security (RLS) to restrict access to a dataset .	December 6, 2020

[New percentile functions](#)

QuickSight now supports variations on percentile calculations, including aggregate functions `percentileCont` and `percentileDisc` and `OVER` functions `percentileContOver` and `percentileDiscOver`. You can use all of these functions in the calculations editor in analysis mode. For more information, see [Functions by category](#).

December 6, 2020

[Dynamic queries, plus dataset editor improvements](#)

QuickSight now enables you to improve performance for visuals based on direct queries by specifying unique keys for joined tables. There are also multiple interface improvements, such as settings for optional autopreviews, zoom in and out of data diagrams, join recommendations, and more. For more information, see [Joining data](#).

December 6, 2020

[Developer portal for embedding](#)

The Dev Portal helps you learn how to use embedding by example in your web site or application. For more information, see [Using the dev portal](#).

November 30, 2020

New data source: Amazon OpenSearch	QuickSight now supports connecting to Amazon OpenSearch. For more information, see Using Amazon OpenSearch with QuickSight .	November 25, 2020
Format metrics on pie charts	QuickSight now supports formatting metrics on pie charts to display values, percentages, or both. For more information, see Customizing data labels .	November 25, 2020
New formatting options for tables	In QuickSight, you can now vertically align and wrap the text for headers in table charts. For more information, see Using tables as visuals .	November 24, 2020
New chart type: Box plots	QuickSight now supports using box plots so you can visualize how your data is distributed across an axis or over time. For more information, see Box plots .	November 24, 2020
Filled maps	QuickSight now supports filled maps so you can visualize your data over a geographical area. For more information, see Filled maps .	November 24, 2020

Favorite folders	In QuickSight, you can now favorite your folders for easy access. For more information, see Organizing assets into folders .	November 24, 2020
Change in display limits	For parameter controls and filters, Amazon QuickSight now displays up to 1,000 sample values. When you have more than that, you use the search box to locate a value. For more information, see Using a control with a parameter .	November 16, 2020
Null rendering	QuickSight now allows null values to be updated to a custom string. For more information, see Changing a field format .	November 9, 2020
New chart type: Waterfall charts	QuickSight now supports using waterfall charts so you can view your data sequentially. For more information, see Waterfall charts .	November 9, 2020

Embedding for everyone	QuickSight Enterprise edition now supports embedding for everyone. When you purchase session packs with capacity pricing, you can enable visitors to use embedded dashboards without registering them as QuickSight users. For more information, see Embedded analytics .	November 9, 2020
Column-level security	QuickSight now supports using column-level security to restrict access to a dataset. For more information, see Using column-level security .	November 9, 2020
Service provider initiated federation	QuickSight Enterprise edition now supports service provider initiated federation, so you can sign on directly to QuickSight using your IAM Identity Center login and password. For more information, see Setting up service provider–Initiated federation with QuickSight enterprise edition .	October 30, 2020
New chart type: Funnel charts	QuickSight now supports using funnel charts so you can display your data in a linear process. For more information, see Funnel charts .	October 29, 2020

Custom sorting	QuickSight now enables you to apply sort options on fields outside the field wells. This feature is available for all charts that support sorting. For more information, see Describing data .	October 29, 2020
New region: AWS GovCloud (US-West)	Amazon QuickSight is now available in AWS GovCloud (US-West) (gov-west-1). For more information, see QuickSight - AWS GovCloud (US) .	October 28, 2020
New data source: Oracle	QuickSight now supports connecting to Oracle. For more information, see Supported data sources .	October 23, 2020
Filter across datasets	In QuickSight, you can now create filters that apply to multiple datasets in an analysis. For more information, see Filtering data .	October 23, 2020
Add field descriptions to datasets	QuickSight now enables you to add column or field descriptions to datasets. This metadata is visible in both datasets and analyses, helping you make your data self-explanatory. For more information, see Describing data .	October 23, 2020

Filter controls on analysis sheets	QuickSight now offers a filter control that you can add to your analysis with a single click. You can put filters beside dashboard visuals and resize them to fit. For more information, see Using filter controls .	October 5, 2020
New data source: Amazon Timestream	QuickSight now supports Amazon Timestream as a data source. For more information, see Using Amazon Timestream data with QuickSight .	October 1, 2020
Export to excel	QuickSight now supports exporting data from pivot tables and table charts to Microsoft Excel (.xlsx) format. For more information, see Exporting data .	September 14, 2020
Duplicate sheets	QuickSight now supports duplicating sheets. For more information, see Working with multiple sheets in an QuickSight analysis .	September 14, 2020
Customize labels of totals and subtotals	In QuickSight, you can now add custom text to display beside totals and subtotals for pivot tables and table charts. For more information, see Table and pivot table formatting options in QuickSight .	September 14, 2020

Skipped row details	When you import data into QuickSight, you now get better error messages and row-by-row information on what caused rows to be skipped during ingestion into SPICE. For more information, see Troubleshooting skipped rows .	September 12, 2020
New configuration option for multivalue lists	QuickSight now supports a configuration option for multivalue lists to start with no values selected. For more information, see Setting up parameters in Amazon QuickSight .	September 11, 2020
Reference lines	QuickSight now supports using reference lines in bar, line, and combo charts. For more information, see Reference lines .	September 3, 2020
New table calculations	QuickSight now supports using window functions <code>firstValue</code> and <code>lastValue</code> in analyses. For more information, see Table calculations .	August 25, 2020
Multivalue default parameters	QuickSight now supports dynamic default values for multi-valued parameters. For more information, see Creating parameter defaults in QuickSight .	August 25, 2020

Last active status for users	QuickSight administrators can now see the last date and time that people accessed QuickSight. For more information, see Viewing user details .	August 25, 2020
Join SQL queries, tables, and files	QuickSight now supports using custom SQL queries in joins with other queries, tables, and files. For more information, see Joining data .	August 25, 2020
Calculated expressions editor	The enhanced editor adds ease-of-use features to help you more easily create calculated expressions in QuickSight. With the new full-screen editor, you can add functions and values directly from the menu. For more information, see Adding a calculated field to an analysis .	August 25, 2020
Personal and shared folders	QuickSight Enterprise edition now offers personal and shared folders to make it easier to discover, organize, share, and explore your available assets. For more information, see Organizing QuickSight assets into folders .	August 5, 2020

[QuickSight support for multitenancy](#)

In Enterprise edition, QuickSight now offers an additional layer of security to support multitenancy. By creating users and groups in a QuickSight namespace, you can isolate them from users and groups in other namespaces. For more information, see [Supporting multitenancy with isolated namespaces](#).

July 23, 2020

[Embed the QuickSight console](#)

In QuickSight Enterprise edition, you can now provide the full dashboard authoring experience of the QuickSight console in your own custom-branded authoring portal. For more information, see [Embedding the QuickSight console](#).

July 23, 2020

[Datasets with 2,000 columns](#)

QuickSight now supports datasets that have up to 2,000 columns. For more information, see [Data source limits](#).

July 23, 2020

[Customized permissions](#)

In QuickSight Enterprise edition, you can now create custom roles to restrict a person from using specific functionality in the QuickSight console, for example to control who can manage data sources and datasets or who can manage or subscribe to email reports. For more information, see [Customizing access to the QuickSight console](#).

July 23, 2020

[Customize the console](#)

In QuickSight Enterprise edition, you can now choose whether to show or hide QuickSight startup samples and videos. You can also create and specify a default theme to add branding to your experience of the QuickSight console. For more information, see [Customizing the QuickSight console](#).

July 23, 2020

[Larger SPICE datasets](#)

QuickSight Enterprise edition now supports SPICE datasets with up to 250 million (250,000,000) rows or 500 GB. For more information, see [Data source limits](#).

July 9, 2020

[Display width settings](#)

In QuickSight, you can now choose between responsive and fixed layouts for your analyses and dashboards. You can also temporarily change your current view from the menu bar. For more information, see [Display settings](#).

July 9, 2020

[Data refresh notifications](#)

QuickSight now supports sending email to notify SPICE dataset owners when data refresh fails. For more information, see [Refreshing data](#).

July 9, 2020

[Color settings for heat maps and tree maps](#)

You can now customize gradient colors for your heat maps and tree maps in QuickSight. For more information, see [Changing colors on heat maps and tree maps](#).

July 9, 2020

[Font choice](#)

You can now customize your theme with a font. Choose from a selection of fonts supported by QuickSight. For more information, see [Using themes in QuickSight](#).

July 8, 2020

[AWS Lake Formation](#)

If you use QuickSight Enterprise edition to query Amazon Athena, you can simplify data access management by using AWS Lake Formation to control permissions to Athena. For more information, see [Authorizing connections through AWS Lake Formation](#).

June 29, 2020

[New chart type](#)

QuickSight now supports using histograms so you can display the distribution of values in your data. For more information, see [Histogram charts](#).

June 12, 2020

[Gauge chart update](#)

You can now add conditional formatting to gauge charts in QuickSight. For more information, see [Adding conditional formatting to visuals](#).

June 12, 2020

[Five new languages](#)

QuickSight now supports Danish, Dutch, Finnish, Norwegian, and Swedish. These languages expand upon the existing 15 languages already available in QuickSight. For more information, see [Choosing a language in Amazon QuickSight](#).

June 12, 2020

[New region: Mumbai](#)

Amazon QuickSight is now available in Asia Pacific (Mumbai) (ap-south-1). For more information, see [AWS Regions, websites, IP address ranges, and endpoints](#).

June 3, 2020

[New SQL editor](#)

QuickSight enhanced the SQL editor that you use to create datasets from SQL queries. The new SQL editor supports syntax highlighting, basic autocomplete, autoindent, and line numbering. Also, you can use the new schema explorer to interactively explore schemas, tables, fields, and data types. For more information, see [Using a SQL query](#).

May 20, 2020

[Amazon SageMaker AI integration now generally available](#)

The integration of SageMaker AI with QuickSight launched in preview just before AWS re:Invent 2019. Now, this feature is generally available in all AWS Regions that are supported by QuickSight. Amazon QuickSight and SageMaker AI together make it faster, easier, and more cost-effective for customers to make use of their machine learning models for visualization and predictions. For more information, see [Amazon QuickSight integration with SageMaker](#).

May 19, 2020

[Update to presto data sources](#)

You can now use nonauthenticated private Presto as a data source. For more information, see [Creating a data source using presto](#).

May 18, 2020

[Rewrite of Amazon QuickSight VPC connection section](#)

Based on your feedback, we completely rewrote this documentation to better explain using Amazon VPC with QuickSight. It includes examples and AWS CLI commands to help you establish the path from Amazon QuickSight to your data source, even if you're using on-premises sources. For more information, see [QuickSight VPC connections](#).

May 18, 2020

[Use logarithmic scale with QuickSight](#)

You can now format visuals using log scale. For more information, see [Changing the visual scale](#).

May 1, 2020

[Severity levels for QuickSight anomaly detection insights](#)

You can now choose the level of severity for anomalies that display in your insight widgets. For more information, see [Using anomaly detection](#).

April 17, 2020

[QuickSight adds a new visual type](#)

The new visual type, stacked area line charts, is similar to area charts except that the stacked values indicate the relationship that each value contributes to the whole. Using stacked area charts, you can display cumulative totals over time while displaying the breakdown by the chosen category. For more information, see [Line charts](#).

April 17, 2020

[QuickSight activates legends by adding menus for quick feature access](#)

QuickSight now supports clicking on legend items to quickly access features like filtering, drilling, chart colors, and custom actions for filtering and URLs. For more information, see [Focusing on visual elements](#) and [Custom actions in QuickSight](#).

April 17, 2020

[Improvements to visual controls in QuickSight](#)

The on-visual menu can now float at the upper-right corner of the visual, moving inside or outside the border of the visual. Plus, you can make your visuals smaller and have more of them, as many as 36 across. For smaller visuals, the new visual controls have more viewing area for charts and insights. To see the new menu, choose one of your visuals. For more information, see [Working with visuals](#).

April 17, 2020

[QuickSight adds cascading filters](#)

You can now create cascading filters by adding them to custom actions in your analyses and dashboards. Each custom filter action can target one or more visuals in the same sheet, creating a cascading filter effect. For more information, see [Custom actions in QuickSight](#).

April 2, 2020

[Modulo operation available in Amazon QuickSight](#)

You can now use modulo operation to find the remainder after dividing one number into another. For more information, see [mod](#).

March 8, 2020

[Images in Amazon QuickSight narratives](#)

You can now insert images from URLs to be rendered as part of a narrative. Images can be resized, conditionally shown within an IF block, and also hyperlinked with the URL feature. For more information, see [Expression editor screen and menus](#).

March 8, 2020

[Amazon QuickSight adds minimum and maximum dates](#)

You can now use minimum and maximum date aggregations in tables and pivot tables. For more information, see [max](#) and [min](#).

March 8, 2020

[Now you can customize minimum delta values to enhance QuickSight ML-powered anomaly detection](#)

Now you can control the minimum delta value that QuickSight uses to detect anomalies (also known as outliers). You can change this setting when you are configuring or exploring anomalies. For more information, see [Detecting outliers with ML-powered anomaly detection](#).

February 4, 2020

[Improvements in the QuickSight narrative expression editor](#)

QuickSight now provides a much more spacious interface for editing narrative expressions. The interface now has instant feedback in a preview of your evaluated narrative before you save any changes. Plus, narratives now support both static and dynamic URLs in your custom narratives for links to searches, apps, other dashboards, other sheets in the same dashboard, and more. For more information, see [Using autonarratives](#).

February 4, 2020

[SageMaker AI integration now available in more AWS Regions](#)

The public preview for integrating SageMaker AI with QuickSight is now available in all of the AWS Regions that are supported by QuickSight. For more information, see [Amazon QuickSight integration with SageMaker](#).

January 10, 2020

[QuickSight now enables you to choose an Amazon Athena workgroup](#)

By choosing an Athena workgroup, you can better manage your Athena datasets. Using this option, you can also allocate Athena costs to the workgroup that QuickSight is using, for better cost reporting. For more information, see [Creating a dataset using Amazon Athena](#).

December 10, 2019

[Private VPC support for presto in QuickSight](#)

Now Presto can use a VPC connection. For more information, see [Supported data sources for VPC](#).

December 10, 2019

[New mathematical functions in QuickSight](#)

Amazon QuickSight now supports functions for base 10 logarithms (log), natural logarithm (ln), absolute value (abs), square root (sqrt), and base of natural log e raised to the power of (exp). For more information, see [Functions and operators](#).

December 10, 2019

[New LAA functions in QuickSight](#)

Now you can use rank, denseRank , and percentileRank for level-aware aggregations. For more information, see [Using level-aware aggregations](#).

December 10, 2019

[Integrate SageMaker AI ML models into your QuickSight analyses and dashboards](#)

In this public preview, QuickSight launches the ability to integrate machine learning (ML) models created and trained in SageMaker AI. This new feature makes it easier to augment your business data with ML predictions. Add your data scientists' prebuilt inferences and predictions to your analysts' dashboards, and let the collaboration and decision-making begin. For more information, see [Amazon QuickSight integration with SageMaker](#).

November 26, 2019

[New formatting options for visuals in QuickSight](#)

Amazon QuickSight launches more formatting for visuals and visuals titles, including font sizes from extra small to extra large. In key performance indicators (KPIs), you can set font sizes for primary and comparison values. In pivot tables and tables, you can see font sizes for table headers, cells, totals, and subtotals. Now your visuals render better even when smaller and with less padding between charts. For combo charts, you can now synchronize the Y-axes for both bars and lines into a single axis. For more information, see [Formatting a visual](#).

November 23, 2019

[QuickSight supports the now function in SPICE](#)

For more information, see [now](#).

November 22, 2019

[QuickSight now supports seconds in SPICE](#)

Seconds are now supported in SPICE datasets. Datetime fields are no longer truncated to minutes. Now you can use the SS option for the period parameter in date functions including `addDateTime`, `dateDiff`, `extract`, and `truncDate`. For more information, see [Date functions](#).

November 22, 2019

[QuickSight now has themes](#)

You can now create a collection of themes, and apply a theme to an analysis and all its dashboards. For more information, see [Using themes in QuickSight](#).

November 22, 2019

[New QuickSight API operations](#)

Amazon QuickSight launches new API operations to programmatically manage your data, dashboards, and fine-grained access control capabilities linked with AWS Identity and Access Management (IAM). With new data API operations, you can create, update, and delete Amazon QuickSight data sources and datasets programmatically. You can also manage data refreshes on your SPICE datasets with API operations. Amazon QuickSight also introduces templates, which store the visual configuration and data schemas required for a dashboard. You can transfer templates across accounts or use them to instantiate dashboards with the same visual presentation but different data. For more information, see the [Amazon QuickSight API reference](#).

November 22, 2019

[QuickSight now supports conditional formatting for tables, pivot tables, and key performance indicators \(KPIs\).](#)

For tables and pivot tables, you can set multiple conditions for fields or supported aggregations, along with format options to apply to a target cell. For KPIs, you can format the primary value based on conditions that are applied to any dimension in the dataset. The conditional formatting options now supported are text color, background color, and placement of supported icons. You can use icons from the provided set, or you can use Unicode icons instead. For more information, see [Adding conditional formatting to visuals](#).

November 18, 2019

[View history for SPICE ingestion on your Amazon QuickSight datasets](#)

You can now view the ingestion history for SPICE datasets in Amazon QuickSight. See information like when the latest ingestion started, how long it took, and what its status is. For more information, see [View SPICE ingestion history](#).

November 7, 2019

[Add your own functionality to visuals with Amazon QuickSight actions](#)

Amazon QuickSight enables you to add to the basic functionality for visuals by creating your own custom actions for filtering or opening URLs. For more information, see [Custom actions in Amazon QuickSight](#).

November 7, 2019

[Amazon QuickSight now supports dynamic sheet titles](#)

You can now use parameters in sheet titles to make the context of a dashboard clearer to the reader. For more information, see [Formatting a visual](#).

November 6, 2019

[Use Amazon QuickSight to join data from multiple data sources](#)

Amazon QuickSight now supports creating datasets that join multiple data sources. For more information, see [Joining data](#).

November 5, 2019

[Amazon QuickSight mobile iOS update and a new Android app](#)

Get access to insights from your data on the fly. Download the updated iOS app or the new Android app. You can browse, favorite, and interact with your dashboards and explore your data with drilldowns and filters. You can stay ahead of the curve by using forecasting. You can get email alerts when unexpected changes happen in your data, and you can share those insights with colleagues. For more information, see [Amazon QuickSight mobile](#).

November 5, 2019

[Amazon QuickSight supports printing](#)

Now you can print a dashboard or an analysis. For more information, see [Printing](#).

October 17, 2019

[Amazon QuickSight supports a new AWS Region](#)

Amazon QuickSight is now available in Asia Pacific (Seoul) (ap-northeast-2). For more information, see [AWS Regions, websites, IP address ranges, and endpoints](#).

October 17, 2019

[Amazon QuickSight now supports visual-level formatting](#)

Now, any formatting you apply from the field wells is applied only to the selected visual. For more information, see [Formatting a visual](#).

October 17, 2019

Transpose tables	In table visuals, you can now transpose columns and rows. For more information, see Tables .	October 10, 2019
Shared data sources	Share data sources with other users and groups. For more information, see Shared data sources .	October 10, 2019
SPICE enhancements	SPICE dashboards now support new wildcard filters on dimensions. You can now filter your data using any of the available wildcard filters: 'contains', 'starts with', 'end with', 'equals'. Additionally, SPICE dashboards now support new string functions (toString and parseDecimal) and two new date functions (parseDate and formatDate). For more information, see Calculated fields .	October 10, 2019
New publishing options for dashboards	Control your user experience with more publishing options for your dashboards. These options include ability to toggle the filter pane, tooltips, drill up/down, and more. For more information, see Publish a dashboard .	October 10, 2019

New median function	Amazon QuickSight supports median as a new aggregation. For more information, see median .	October 10, 2019
New filter functionality	Amazon QuickSight now supports additional relative date filter options. For more information, see Date filters .	October 10, 2019
New aggregations for fields	You can now use Nth percentile, median, standard deviation, and variance in field wells, filters, and ML Insights. For more information, see Field aggregations .	October 10, 2019
Filter for context	You can now filter text using the following wildcard comparison types: contains, does not contain, begins with, and ends with. For more information, see Text filters .	October 10, 2019
Use the power operator (^) in SPICE	QuickSight now supports using the power operator (^) on SPICE datasets. You can use the power operator with any numeric field, with any valid exponent. For example, "revenue ^ 0.5 " computes the square root of the revenue field. For more information, see Arithmetic and comparison operators .	September 11, 2019

[Use string functions in SPICE](#)

Now you can use string functions for calculated fields in your analyses based on SPICE datasets. For more information, see [Functions by category](#).

September 11, 2019

[Use level-aware aggregations to control your calculations](#)

QuickSight now supports level-aware aggregation calculations. Using this feature, you can perform aggregations at prefilter and preaggregation levels, before aggregations that happen in the display. This helps you explore more advanced query-building strategies than ever before. For more information, see [Level-aware aggregations](#).

September 11, 2019

[Sort your anomalies](#)

Now you can choose how to sort anomalies on the insight widget and in the anomaly exploration page. Choosing how to prioritize anomalies can help you identify the anomalies that are the most important to you. For more information, see [Using ML-powered anomaly detection](#).

September 11, 2019

[Put 100 million rows into a SPICE dataset](#)

We changed the SPICE dataset limit from 25GB to 100 million rows for Enterprise edition (and 25 million rows for Standard edition). For more information, see [SPICE data source limits](#).

September 11, 2019

[Rename your dashboard, keep your settings](#)

Amazon QuickSight now supports renaming your published dashboards in place. You can change the name and still keep your settings and subscribers. For more information, see [Publishing a dashboards](#).

September 6, 2019

[New visual type: Word clouds](#)

Amazon QuickSight supports word cloud visuals that are based on different aggregations over any dimension in your dataset. You can drill down, focus on specific data points, and apply specific colors to data points on the word cloud. For more information, see [Word clouds](#).

September 5, 2019

[Filter out the last \$n\$ time periods](#)

Amazon QuickSight lets you exclude a specific number and type of time periods from a time range (after) filter. For more information, see [Adding a date filter](#).

September 5, 2019

[Amazon QuickSight new feature to mark favorites](#)

Mark your favorite dashboards and analyses so you can come back to them quickly. For more information, see [Using the QuickSight start page](#).

September 5, 2019

[Amazon QuickSight Enterprise edition adds anomaly alerts](#)

Your readers can subscribe themselves to anomaly alerts on dashboards to get the latest ML-powered anomaly detection and contribution analysis by email. For more information, see [Using data dashboards](#).

September 5, 2019

[QuickSight adds custom colors for charts](#)

With the QuickSight enhanced color picker, you can choose custom colors for charts where color customization is supported. The enhanced color picker also retains the eight most recently used custom colors for easy selection across multiple charts. For more information, see [Changing visual colors](#).

August 1, 2019

[Detect anomalies with less data](#)

In Amazon QuickSight Enterprise edition, we lowered the required minimum data points to 15 for training anomaly detection. For more information, see [Dataset requirements for using ML insights with QuickSight](#).

August 1, 2019

[Amazon QuickSight supports new AWS Regions](#)

Amazon QuickSight is available in Europe (London) and Europe (Frankfurt). Also, you can now connect to Active Directory in any AWS Region supported by Amazon QuickSight, except Asia Pacific (Singapore) and Asia Pacific (Sydney). For more information, see [AWS Regions and IP address ranges](#).

August 1, 2019

[Embedding for Active Directory](#)

In Amazon QuickSight Enterprise edition, you can now embed dashboards for users authenticated through Active Directory. For more information, see [Embedding dashboards](#).

July 11, 2019

[Granular access control for using AWS services from Amazon QuickSight](#)

In Amazon QuickSight Enterprise edition, you can scope down your security policies to allow specific people access to specific AWS resources. You can attach and detach IAM policies in the Amazon QuickSight interface, simplifying finer control over who can access your AWS data sources. For more information, see [Controlling access to AWS resources](#).

June 14, 2019

[Tabs for dashboards and analyses](#)

Amazon QuickSight now supports multiple sheets inside of analyses and dashboards. You can add multiple tabs to your dashboards, utilize URL actions and on-screen controls to simplify navigation, and filter across all of your sheets. For more information, see [Multiple sheets](#).

June 11, 2019

[New visual types](#)

You can now use Amazon QuickSight to chart your data on a customizable gauge and donut charts. For more information, see [Visual types](#).

June 11, 2019

[New functions](#)

Amazon QuickSight now supports functions Variance and Standard Deviations as both aggregation and table calculation. You can create Sample and Population variants for both the functions. You can create these calculations through the calculations editor in analysis mode. [Functions by category](#)

June 11, 2019

[Hide/Show columns in table visualizations](#)

In Amazon QuickSight, you can hide or show columns in visualizations that use the *table* visual type. For more information, see [Customizing a visual](#).

June 11, 2019

Conditional string functions on SPICE	We added support for <code>ifelse</code> in analyses that use SPICE. For more information, see ifelse	June 11, 2019
Choose period granularity for date differences	We added an optional parameter to <code>dateDiff</code> so you can show differences between dates in the period you choose (years, months, and so on). For more information, see dateDiff .	June 11, 2019
Duplicate visuals	In Amazon QuickSight, you can now duplicate visuals. For more information, see Duplicating a visual .	May 21, 2019
Angled data labels	In Amazon QuickSight, data labels that are too long are now angled by default on vertical bar, combo, and line charts. For more information, see Customizing data labels .	May 21, 2019
New languages added	Amazon QuickSight is now available in 10 languages: English, German, Spanish, French, Portuguese, Italian, Japanese, Korean, Simplified Chinese, and Traditional Chinese. For more information, see Choosing a language in Amazon QuickSight .	April 8, 2019

[New aggregation function](#)

Amazon QuickSight supports aggregating by percentile. This function helps you understand the distribution of your data. For more information, see [percentile](#).

April 8, 2019

[Customize how many data points to display](#)

You can now format your visuals to display a custom number of data points or groups before showing the "other" category. This feature is available for bar charts, combo charts, line charts, pie charts, heat maps, and tree maps. For more information, see [Customizing a visual](#).

April 8, 2019

[ML-powered forecasting](#)

With forecasting powered by machine learning and what-if analyses in Amazon QuickSight, nontechnical users can now easily forecast their key business metrics. No ML expertise or Microsoft Excel data modeling is required. The built-in ML algorithm in Amazon QuickSight is designed to handle complex real-world scenarios. Amazon QuickSight uses ML to provide more reliable forecasts than traditional means. For more information, see [Forecasts and what-ifs](#).

March 14, 2019

[ML-powered anomaly detection](#)

Amazon QuickSight uses proven Amazon technology to continuously run ML-powered anomaly detection on millions of metrics and billions of data points. This anomaly detection enables you to get deep insights that are often buried in the aggregates, not visible in plain sight, and not scalable with manual analysis.

With ML-powered anomaly detection, there's no need for manual analysis, custom development, or ML domain expertise. For more information, see [Anomaly detection](#).

March 14, 2019

[Automatic narratives](#)

Automatic narratives provide key insights in everyday language, embedded contextually in your dashboard, saving hours on manual analysis. With automatic narratives, Amazon QuickSight interprets the charts and tables in your dashboard and provides a number of suggested insights in natural language. Depending on the shape and form of your data, you might get different suggestions. For example, you might see what the day-over-day changes look like, what was the highest sales date, or what the growth rate is. Or you might see what the forecast looks like for the next seven days. As the author of the dashboard, you can customize the computations and business language for your needs. You can use automatic narratives to effectively tell the story of your data in plain language. For more information, see [Narrative insights](#).

March 14, 2019

Previous updates

The following table describes the important changes in each release of the *Amazon Quick User Guide* before March 4, 2019.

Change	Description	Date changed
New features	Using custom window functions, you can calculate any aggregation for a defined window at the point of selection. You can define the window interval before and after any point in time. You can also do calculations for that window using the sum, average, minimum, maximum, and count functions. Doing this enables a moving window aggregation as you progress through your data points. For more information, see Table calculation functions. Send email reports with data tailored to each of your users and groups. You can now create email reports for datasets that use row-level security. Amazon QuickSight generates a custom email snapshot for each user or group based on their data permission that is defined in the dashboard. RLS for email reports works for both scheduled and ad hoc emails. For more information, see Scheduling and sending QuickSight reports by email. Amazon QuickSight now supports unbounded cardinality. That means you can have more than 10,000 values in your control or filter. For more information, see Using a control with a parameter in Amazon QuickSight. You can create a datetime parameter that has no static default value. Filters on these parameters become active after you choose a value. For more information, see Adding date filters.	March 4, 2019

Change	Description	Date changed
New features	There are new table calculations available: • percentileOver • runningAvg • runningCount • runningMax • runningMin	February 7, 2019
New features	Amazon QuickSight supports emailing reports that use row-level security, using controls with unbounded cardinality, and creating datetime parameters with no default values. For more information, see the following: • Scheduling and sending QuickSight reports by email • Using a control with a parameter in Amazon QuickSight • Adding date filters	January 22, 2019

Change	Description	Date changed
New features	We added conditional aggregations, including <code>sumIf</code>, <code>countIf</code>, <code>minIf</code>, <code>maxIf</code>, <code>avgIf</code>, and <code>distinct_countIf</code> . Pivot tables now support infinite scrolling through millions of rows. You can add up to 20 fields for columns, and 20 fields for rows. Plus, you can add subtotals and totals to rows and columns. For more information, see the following: • Aggregate functions • Using pivot tables	January 10, 2019
New feature	We enhanced the join editor to increase usability and functionality. You can now add tables from one or more schemas on the same data source, or add the same table twice. For more information, see Joining data.	January 3, 2019
New SDK	You can embed dashboards and manage users or groups by using the Amazon QuickSight SDKs. For more information, see Developing with QuickSight.	November 27, 2018
New features	You can use groups with row-level security (RLS), and you can add cascading controls for parameters. For more information, see the following: • Using row-level security with user-based rules to restrict access to a dataset • Using a control with a parameter in Amazon QuickSight	November 20, 2018

Change	Description	Date changed
Function renamed	We renamed the percentile function to make it more intuitive. Its new name is <code>percentileRank</code> . For more information, see percentileRank .	November 12, 2018
New features	Use a top and bottom filter to show the top or bottom <code>n</code> for the field you choose, based on values in another field. For example, you could choose to show the top five sales people based on revenue. For more information, see Adding text filters. Use cascading controls to limit the values displayed in the controls, so they only show values that are relevant to what is selected in other controls. For more information, see Setting up parameters in Amazon QuickSight.	November 1, 2018
New feature	Use JSON native data types with <code>parseJson</code> . For more information see, parseJson .	October 30, 2018
New features	Use date functions to find out what quarter a date is in: • addDateTime • Extract • truncDate Add parameters to URLs. For more information see, Using parameters in a URL.	September 10, 2018
New features	You can sort strings in SPICE datasets. For more information on these features, see Sorting visual data in Amazon QuickSight	August 20, 2018

Change	Description	Date changed
New features	You can schedule emailed reports, and add data labels to your visuals. For more information on these features, see the following sections: • Scheduling and sending QuickSight reports by email • Subscribing to email reports in QuickSight • Data labels on visual types in QuickSight	August 15, 2018
New features	You can create table calculations, using aggregated measures to discover how dimensions influence measures or each other. Also, you can visualize time data at granularities as low as one minute. For more information on these features, see the following sections: • Table calculations • Changing date field granularity	August 8, 2018

Change	Description	Date changed
New features	You can replace datasets, customize labels, and format dimensions that are aggregated with count and count distinct. Also, new visuals start out smaller in size. For more information on these features, see the following sections: • Replacing datasets • Labels on visual types in QuickSight • Customizing a field format • Visualizing data in QuickSight	June 21, 2018

Change	Description	Date changed
New features and a new AWS Region	You can upgrade your Amazon QuickSight subscription from Standard edition to Enterprise edition. In Enterprise edition, Amazon QuickSight supports usage-based pricing for users in the reader role, sharing dashboards with all users in the reader role, and hourly refresh of datasets. Amazon QuickSight also supports private connections to data in a VPC with a private subnet in Enterprise edition. In both editions, Amazon QuickSight supports parameters with on-sheet controls, dashboard co-ownership, custom URL actions, and 25-GB SPICE datasets. Also, Amazon QuickSight is available in Asia Pacific (Tokyo). For more information on these features, see the following sections: • Upgrading your Amazon QuickSight subscription from Standard edition to Enterprise edition • Self-provisioning an Amazon QuickSight read-only user • Inviting users to access Amazon QuickSight • Sharing QuickSight dashboards • Connecting to a VPC with QuickSight • Refreshing a dataset on a schedule • Parameters in Amazon QuickSight •	May 30, 2018

Change	Description	Date changed
	Exploring interactive dashboards in QuickSight (newly updated for read-only users) Using custom actions for filtering and navigating Data source quotas AWS Regions, websites, IP address ranges, and endpoints	
New feature	Amazon QuickSight supports quick sorting from axis labels, duplicating datasets, and showing or hiding totals on tabular reports. Also, you can add custom SQL to a dataset earlier in the creation workflow. For more information on these features, see the following sections: Sorting visual data in Amazon QuickSight Duplicating datasets Totals and subtotals Creating a basic SQL query	May 25, 2018
New features	You can use Amazon QuickSight to connect to Software as a Service (SaaS) providers. For more information, see Software as a service (SaaS) data. You can import JSON files to Amazon QuickSight. For more information, see JSON data. You can also parse JSON fields in a CSV file. For more information, see parseJson.	April 9, 2018

Change	Description	Date changed
New feature	You can use Amazon QuickSight with Amazon S3 files that are in a different AWS account. For more information, see Datasets using S3 files in another AWS account .	November 20, 2017
New visual types	You can create visuals using maps. You can also view tabular data in a visual. For more information on preparing geospatial data for use in a visual, see Adding geospatial data . For more information on creating a geospatial visual, see Creating maps and geospatial charts . For more information on using tabular reports, see Using tables as visuals .	November 20, 2017
New features	Amazon QuickSight can support 1000 columns in a dataset. For more information, see Data source quotas . Calculated fields are supported in SPICE datasets. For more information, see Adding calculated fields . Also, high cardinality values, which often display as a long tail on a visual, are placed into a category called other . For more information, see Visual types in Amazon QuickSight .	November 20, 2017
New feature	In Enterprise edition, you can restrict access to a dataset by adding row-level security. To learn more, see Using row-level security with user-based rules to restrict access to a dataset .	October 20, 2017
New visual type	You can create visuals using combo charts. To learn more about combo charts, see Using combo charts .	October 20, 2017

Change	Description	Date changed
New features	Amazon QuickSight supports creating custom aggregations for calculated fields in analyses, custom date formats, and copies of dashboards. For more information on aggregating calculated fields, see Aggregate functions . For more information on using unsupported dates by creating a custom date format, see Using unsupported or custom dates . For more information on duplicating dashboards, see Copying an QuickSight dashboard .	September 25, 2017
New feature	You can combine multiple filters using the And/Or operators. To learn more about filter groups, see Adding filter conditions (group filters) with AND and OR operators .	August 31, 2017
New data source	Amazon QuickSight supports Amazon S3 Analytics.	August 31, 2017
New Features	Amazon QuickSight supports importing ZIP files from Amazon S3. There is also a new search feature, to simplify finding analyses, datasets, and dashboards. For more information on the search feature, see Using the QuickSight console .	August 31, 2017
New AWS Regions	Amazon QuickSight is now available in Asia Pacific (Singapore) and Asia Pacific (Sydney).	August 8, 2017
New feature	Amazon QuickSight supports Snowflake cloud data warehouse.	July 31, 2017
New feature	Amazon QuickSight added a new aggregation: Count Distinct. To learn more, see Changing field aggregation .	July 19, 2017

Change	Description	Date changed
New feature	Amazon QuickSight supports exploring Amazon S3 analytics data from an Amazon QuickSight dashboard that you reach from the AWS Management Console. To learn more, see Exploring your AWS data in Amazon QuickSight .	July 5, 2017
New feature	Amazon QuickSight supports Federated Single Sign-On (IAM Identity Center) in Enterprise Edition. To learn more, see Using external identity federation and single sign-on with Amazon QuickSight .	May 25, 2017
New feature	Amazon QuickSight supports Amazon Redshift Spectrum. To learn more, see Enabling access to Amazon Redshift Spectrum .	May 25, 2017
New feature	Amazon QuickSight supports Federated Single Sign-On (IAM Identity Center) in Standard Edition. To learn more, see Using external identity federation and single sign-on with Amazon QuickSight .	May 25, 2017
New feature	Amazon QuickSight supports just-in-time (JIT) user provisioning through the following policy actions: <code>quicksight:CreateUser</code> and <code>quicksight:CreateAdmin</code> . To learn more, see IAM policy examples for QuickSight .	May 25, 2017
New feature	Amazon QuickSight supports direct connections to Teradata 14.0 and later.	May 25, 2017
New feature	Amazon QuickSight added relative date filters for datasets and visuals. To learn more, see Adding date filters .	May 25, 2017

Change	Description	Date changed
New feature	Amazon QuickSight supports connecting to Apache Spark and Presto. To learn more, see Creating a data source using Apache Spark and Creating a data source using Presto .	May 3, 2017
New feature	Amazon QuickSight supports operational logging with AWS CloudTrail. To learn more, see Logging operations with AWS CloudTrail .	April 28, 2017
New feature	Amazon QuickSight is available in US East (Ohio). To learn more about AWS Regions, see AWS Regions, websites, IP address ranges, and endpoints .	April 11, 2017
New feature	- You can use the AD Connector with Amazon QuickSight. To learn more about managing Enterprise users, see Managing user access inside Amazon QuickSight. - You can add key performance indicators (KPIs) to your visualizations. To learn more, see Using KPIs. - You can import .xlsx files with headers and footers, comments, formatting, filter header, sort, frozen panel/header, hidden columns, groups, and formulas / references. You can also limit your import to a specific range. To learn more about importing ranges, see Choosing file upload settings.	April 6, 2017
New feature	You can now export to a comma-separated value (CSV) format file using Amazon QuickSight. To learn more, see Exporting data from visuals .	March 21, 2017
New feature	You can now schedule data refreshes for SPICE data sets. To learn more, see Refreshing SPICE data .	February 14, 2017

Change	Description	Date changed
New feature	You can now connect to Amazon Athena databases and use them as data sources in Amazon QuickSight. To learn more, see Creating a dataset using Amazon Athena data .	December 22, 2016
New edition	Amazon QuickSight now offers an Enterprise edition as well as a Standard edition. Both editions offer a full set of features for creating and sharing data visualizations, and Enterprise edition additionally offers encryption at rest and Active Directory integration. When you choose to use Enterprise edition, you select a Microsoft Active Directory directory in Directory Service and use that active directory to identify and manage your Amazon QuickSight users and administrators. To learn more, see Different editions of Amazon QuickSight .	December 15, 2016
New guide	This is the first release of <i>QuickSight User Guide</i> .	November 15, 2016

Attributions

Following, you can find out about attributions and third-party licenses for material used in Quick documentation.

** Mapzen Who's On First; version latest -- <https://tinyurl.com/y8p87rwl>

Data from Who's On First. <http://whosonfirst.mapzen.com#License>

CC0 1.0 Universal

CREATIVE COMMONS CORPORATION IS NOT A LAW FIRM AND DOES NOT PROVIDE LEGAL SERVICES. DISTRIBUTION OF THIS DOCUMENT DOES NOT CREATE AN ATTORNEY-CLIENT RELATIONSHIP. CREATIVE COMMONS PROVIDES THIS INFORMATION ON AN "AS-IS" BASIS. CREATIVE COMMONS MAKES NO WARRANTIES REGARDING THE USE OF THIS DOCUMENT OR THE INFORMATION OR WORKS PROVIDED HEREUNDER, AND DISCLAIMS LIABILITY FOR DAMAGES RESULTING FROM THE USE OF THIS DOCUMENT OR THE INFORMATION OR WORKS PROVIDED HEREUNDER.

Statement of Purpose

The laws of most jurisdictions throughout the world automatically confer exclusive Copyright and Related Rights (defined below) upon the creator and subsequent owner(s) (each and all, an "owner") of an original work of authorship and/or a database (each, a "Work").

Certain owners wish to permanently relinquish those rights to a Work for the purpose of contributing to a commons of creative, cultural and scientific works ("Commons") that the public can reliably and without fear of later claims of infringement build upon, modify, incorporate in other works, reuse and redistribute as freely as possible in any form whatsoever and for any purposes, including without limitation commercial purposes. These owners may contribute to the Commons to promote the ideal of a free culture and the further production of creative, cultural and scientific works, or to gain reputation or greater distribution for their Work in part through the use and efforts of others.

For these and/or other purposes and motivations, and without any expectation of additional consideration or compensation, the person associating CC0 with a Work (the "Affirmer"), to the extent that he or she is an owner of Copyright and Related Rights in the Work, voluntarily elects to

apply CC0 to the Work and publicly distribute the Work under its terms, with knowledge of his or her Copyright and Related Rights in the Work and the meaning and intended legal effect of CC0 on those rights.

1. Copyright and Related Rights. A Work made available under CC0 may be protected by copyright and related or neighboring rights ("Copyright and Related Rights"). Copyright and Related Rights include, but are not limited to, the following:

- i. the right to reproduce, adapt, distribute, perform, display, communicate, and translate a Work;
- ii. moral rights retained by the original author(s) and/or performer(s);
- iii. publicity and privacy rights pertaining to a person's image or likeness depicted in a Work;
- iv. rights protecting against unfair competition in regards to a Work, subject to the limitations in paragraph 4(a), below;
- v. rights protecting the extraction, dissemination, use and reuse of data in a Work;
- vi. database rights (such as those arising under Directive 96/9/EC of the European Parliament and of the Council of 11 March 1996 on the legal protection of databases, and under any national implementation thereof, including any amended or successor version of such directive); and
- vii. other similar, equivalent or corresponding rights throughout the world based on applicable law or treaty, and any national implementations thereof.

2. Waiver. To the greatest extent permitted by, but not in contravention of, applicable law, Affirmer hereby overtly, fully, permanently, irrevocably and unconditionally waives, abandons, and surrenders all of Affirmer's Copyright and Related Rights and associated claims and causes of action, whether now known or unknown (including existing as well as future claims and causes of action), in the Work (i) in all territories worldwide, (ii) for the maximum duration provided by applicable law or treaty (including future time extensions), (iii) in any current or future medium and for any number of copies, and (iv) for any purpose whatsoever, including without limitation commercial, advertising or promotional purposes (the "Waiver"). Affirmer makes the Waiver for the benefit of each member of the public at large and to the detriment of Affirmer's heirs and successors, fully intending that such Waiver shall not be subject to revocation, rescission, cancellation, termination, or any other legal or equitable action to disrupt the quiet enjoyment of the Work by the public as contemplated by Affirmer's express Statement of Purpose.

3. Public License Fallback. Should any part of the Waiver for any reason be judged legally invalid or ineffective under applicable law, then the Waiver shall be preserved to the maximum extent

permitted taking into account Affirmer's express Statement of Purpose. In addition, to the extent the Waiver is so judged Affirmer hereby grants to each affected person a royalty-free, non transferable, non sublicensable, non exclusive, irrevocable and unconditional license to exercise Affirmer's Copyright and Related Rights in the Work (i) in all territories worldwide, (ii) for the maximum duration provided by applicable law or treaty (including future time extensions), (iii) in any current or future medium and for any number of copies, and (iv) for any purpose whatsoever, including without limitation commercial, advertising or promotional purposes (the "License"). The License shall be deemed effective as of the date CC0 was applied by Affirmer to the Work. Should any part of the License for any reason be judged legally invalid or ineffective under applicable law, such partial invalidity or ineffectiveness shall not invalidate the remainder of the License, and in such case Affirmer hereby affirms that he or she will not (i) exercise any of his or her remaining Copyright and Related Rights in the Work or (ii) assert any associated claims and causes of action with respect to the Work, in either case contrary to Affirmer's express Statement of Purpose.

4. Limitations and Disclaimers.

a. No trademark or patent rights held by Affirmer are waived, abandoned, surrendered, licensed or otherwise affected by this document.

b. Affirmer offers the Work as-is and makes no representations or warranties of any kind concerning the Work, express, implied, statutory or otherwise, including without limitation warranties of title, merchantability, fitness for a particular purpose, non infringement, or the absence of latent or other defects, accuracy, or the present or absence of errors, whether or not discoverable, all to the greatest extent permissible under applicable law.

c. Affirmer disclaims responsibility for clearing rights of other persons that may apply to the Work or any use thereof, including without limitation any person's Copyright and Related Rights in the Work. Further, Affirmer disclaims responsibility for obtaining any necessary consents, permissions or other rights required for any use of the Work.

d. Affirmer understands and acknowledges that Creative Commons is not a party to this document and has no duty or obligation with respect to this CC0 or use of the Work.

** Quattroshapes; version latest -- <https://tinyurl.com/y72sd8u8>

Includes data from foursquare quattroshapes. <https://github.com/foursquare/quattroshapes/>.

This map database contains open data from government and other sources. Consider including the following copyright statements when using the data.

Europe

EuroGeoGraphics data copyright is held by European National Mapping Agencies.

Austria © Bundesamt für Eich- und Vermessungswesen

Belgium © Institut Géographique National – Belgique

Bulgaria © Cadastre Agency, MRDPW

Croatia © Drzavna Geodetska Uprava

Cyprus © Lands and Surveys, Survey and Cartogr. Br.

Czech Republic © Český úřad zeměměřický a katastrální

Denmark © Kort og Matrikelstyrelsen

Estonia © Maaamet

Faroe Islands © Kort og Matrikelstyrelsen

Finland © Maanmittauslaitos

France © Institut National de l'Information Géographique et Forestière – France

Georgia © National Agency of Public Registry

Germany © Bundesamt für Kartographie und Geodäsie

Great Britain © Ordnance Survey

Greece © Hellenic Military Geographical Service

Greenland © Kort og Matrikelstyrelsen

Hungary © Földmérési és Távérzékelési Intézet

Iceland © Landmælingar Íslands

Ireland © Ordnance Survey Ireland

Italy © Istituto Geografico Militare Italiano

Latvia © Latvijas Republikas Valsts zemes dienests

Liechtenstein © Bundesamt für Landestopographie (Switzerland)

Lithuania © Nacionalinė žemės tarnyba

Luxembourg © Administration du Cadastre et de la Topographie

Malta © Malta Environment and Planning Authority (MEPA)

Moldova © State Agency for Land Relations and Cadastre

Netherlands © Topografische Dienst Nederland

Northern Ireland © Ordnance Survey of Northern Ireland

Norway © Statens Kartverk

Poland © Główny Urząd Geodezji i Kartografii

Portugal © Instituto Geográfico Português

Rep. of Slovakia © Geodetick a kartografick ústav

Romania © CNGCFT (National Center of Geodesie, Cartography, Fotogrametry and Remote Sensing)

Serbia © Republički geodetski zavod

Slovenia © Geodetska Uprava Republike Slovenije

Spain © Centro Nacional de Informacion Geografica – Instituto Geográfico Nacional

Sweden © Lantmäteriet

Switzerland © Bundesamt für Landestopographie

Ukraine © Research Institute of Geodesy and Cartography

Additional European data

United Kingdom: Contains Ordnance Survey data © Crown copyright and database right [2012]

Netherlands: Kadaster

Spain: Instituto Geográfico Nacional

France: Institut Géographique National

Switzerland: swisstopo

Europe-wide: European Environment Agency (EEA) urban morphological zones 2006

Americas

United States: US Census Bureau (Census 2010 geography files).

Canada: © Department of Natural Resources Canada. All rights reserved., Statistics Canada, and BC Stats

Brazil: IBGE

Mexico: INEGI

Chile: Global Map of Chile © International Steering Committee for Global Mapping / Instituto Geografico Militar de Chile

Asia

Indonesia: Global Map of Indonesia @ ISCGM/Indonesia

Australian: Geoscience Australia and Australian Bureau of Statistics

New Zealand: Land Information New Zealand

South Africa: Global Map of the Republic of South Africa © ISCGM/State Copyright, Department of Rural Development and Land Reform, Chief Directorate of National Geo-spatial Information, Website: <http://www.ngi.gov.za/>

Palestine: "Global Map of PALESTINE © ISCGM/Palestinian National Authority, Ministry of Planning. Geographic Center and Technical Support www.mop.gov.ps

Ordnance Survey

Contains OS data © Crown copyright and database right 2016

Contains Royal Mail data © Royal Mail copyright and Database right 2016

Contains National Statistics data © Crown copyright and database right 2016

Attribution 2.0

CREATIVE COMMONS CORPORATION IS NOT A LAW FIRM AND DOES NOT PROVIDE LEGAL SERVICES. DISTRIBUTION OF THIS LICENSE DOES NOT CREATE AN ATTORNEY-CLIENT RELATIONSHIP. CREATIVE COMMONS PROVIDES THIS INFORMATION ON AN "AS-IS" BASIS. CREATIVE COMMONS MAKES NO WARRANTIES REGARDING THE INFORMATION PROVIDED, AND DISCLAIMS LIABILITY FOR DAMAGES RESULTING FROM ITS USE.

License

THE WORK (AS DEFINED BELOW) IS PROVIDED UNDER THE TERMS OF THIS CREATIVE COMMONS PUBLIC LICENSE ("CCPL" OR "LICENSE"). THE WORK IS PROTECTED BY COPYRIGHT AND/OR OTHER APPLICABLE LAW. ANY USE OF THE WORK OTHER THAN AS AUTHORIZED UNDER THIS LICENSE OR COPYRIGHT LAW IS PROHIBITED.

BY EXERCISING ANY RIGHTS TO THE WORK PROVIDED HERE, YOU ACCEPT AND AGREE TO BE BOUND BY THE TERMS OF THIS LICENSE. THE LICENSOR GRANTS YOU THE RIGHTS CONTAINED HERE IN CONSIDERATION OF YOUR ACCEPTANCE OF SUCH TERMS AND CONDITIONS.

1. Definitions

a. "Collective Work" means a work, such as a periodical issue, anthology or encyclopedia, in which the Work in its entirety in unmodified form, along with a number of other contributions, constituting separate and independent works in themselves, are assembled into a collective whole. A work that constitutes a Collective Work will not be considered a Derivative Work (as defined below) for the purposes of this License.

b. "Derivative Work" means a work based upon the Work or upon the Work and other pre-existing works, such as a translation, musical arrangement, dramatization, fictionalization, motion picture version, sound recording, art reproduction, abridgment, condensation, or any other form in which the Work may be recast, transformed, or adapted, except that a work that constitutes a Collective Work will not be considered a Derivative Work for the purpose of this License. For the avoidance of doubt, where the Work is a musical composition or sound recording, the synchronization of the

Work in timed-relation with a moving image ("synching") will be considered a Derivative Work for the purpose of this License.

c. "Licensor" means the individual or entity that offers the Work under the terms of this License.

d. "Original Author" means the individual or entity who created the Work.

e. "Work" means the copyrightable work of authorship offered under the terms of this License.

f. "You" means an individual or entity exercising rights under this License who has not previously violated the terms of this License with respect to the Work, or who has received express permission from the Licensor to exercise rights under this License despite a previous violation.

2. Fair Use Rights. Nothing in this license is intended to reduce, limit, or restrict any rights arising from fair use, first sale or other limitations on the exclusive rights of the copyright owner under copyright law or other applicable laws.

3. License Grant. Subject to the terms and conditions of this License, Licensor hereby grants You a worldwide, royalty-free, non-exclusive, perpetual (for the duration of the applicable copyright) license to exercise the rights in the Work as stated below:

a. to reproduce the Work, to incorporate the Work into one or more Collective Works, and to reproduce the Work as incorporated in the Collective Works;

b. to create and reproduce Derivative Works;

c. to distribute copies or phonorecords of, display publicly, perform publicly, and perform publicly by means of a digital audio transmission the Work including as incorporated in Collective Works;

d. to distribute copies or phonorecords of, display publicly, perform publicly, and perform publicly by means of a digital audio transmission Derivative Works.

e. For the avoidance of doubt, where the work is a musical composition:

i. Performance Royalties Under Blanket Licenses. Licensor waives the exclusive right to collect, whether individually or via a performance rights society (e.g. ASCAP, BMI, SESAC), royalties for the public performance or public digital performance (e.g. webcast) of the Work.

ii. Mechanical Rights and Statutory Royalties. Licensor waives the exclusive right to collect, whether individually or via a music rights agency or designated agent (e.g. Harry Fox Agency), royalties for any phonorecord You create from the Work ("cover version") and distribute, subject to the

compulsory license created by 17 USC Section 115 of the US Copyright Act (or the equivalent in other jurisdictions).

f. Webcasting Rights and Statutory Royalties. For the avoidance of doubt, where the Work is a sound recording, Licensors waive the exclusive right to collect, whether individually or via a performance-rights society (e.g. SoundExchange), royalties for the public digital performance (e.g. webcast) of the Work, subject to the compulsory license created by 17 USC Section 114 of the US Copyright Act (or the equivalent in other jurisdictions).

The above rights may be exercised in all media and formats whether now known or hereafter devised. The above rights include the right to make such modifications as are technically necessary to exercise the rights in other media and formats. All rights not expressly granted by Licensors are hereby reserved.

4. Restrictions. The license granted in Section 3 above is expressly made subject to and limited by the following restrictions:

a. You may distribute, publicly display, publicly perform, or publicly digitally perform the Work only under the terms of this License, and You must include a copy of, or the Uniform Resource Identifier for, this License with every copy or phonorecord of the Work You distribute, publicly display, publicly perform, or publicly digitally perform. You may not offer or impose any terms on the Work that alter or restrict the terms of this License or the recipients' exercise of the rights granted hereunder. You may not sublicense the Work. You must keep intact all notices that refer to this License and to the disclaimer of warranties. You may not distribute, publicly display, publicly perform, or publicly digitally perform the Work with any technological measures that control access or use of the Work in a manner inconsistent with the terms of this License Agreement. The above applies to the Work as incorporated in a Collective Work, but this does not require the Collective Work apart from the Work itself to be made subject to the terms of this License. If You create a Collective Work, upon notice from any Licensors You must, to the extent practicable, remove from the Collective Work any reference to such Licensors or the Original Author, as requested. If You create a Derivative Work, upon notice from any Licensors You must, to the extent practicable, remove from the Derivative Work any reference to such Licensors or the Original Author, as requested.

b. If you distribute, publicly display, publicly perform, or publicly digitally perform the Work or any Derivative Works or Collective Works, You must keep intact all copyright notices for the Work and give the Original Author credit reasonable to the medium or means You are utilizing by conveying the name (or pseudonym if applicable) of the Original Author if supplied; the title of the Work if supplied; to the extent reasonably practicable, the Uniform Resource Identifier, if any, that Licensors

specifies to be associated with the Work, unless such URI does not refer to the copyright notice or licensing information for the Work; and in the case of a Derivative Work, a credit identifying the use of the Work in the Derivative Work (e.g., "French translation of the Work by Original Author," or "Screenplay based on original Work by Original Author"). Such credit may be implemented in any reasonable manner; provided, however, that in the case of a Derivative Work or Collective Work, at a minimum such credit will appear where any other comparable authorship credit appears and in a manner at least as prominent as such other comparable authorship credit.

5. Representations, Warranties and Disclaimer

UNLESS OTHERWISE MUTUALLY AGREED TO BY THE PARTIES IN WRITING, LICENSOR OFFERS THE WORK AS-IS AND MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND CONCERNING THE WORK, EXPRESS, IMPLIED, STATUTORY OR OTHERWISE, INCLUDING, WITHOUT LIMITATION, WARRANTIES OF TITLE, MERCHANTIBILITY, FITNESS FOR A PARTICULAR PURPOSE, NONINFRINGEMENT, OR THE ABSENCE OF LATENT OR OTHER DEFECTS, ACCURACY, OR THE PRESENCE OF ABSENCE OF ERRORS, WHETHER OR NOT DISCOVERABLE. SOME JURISDICTIONS DO NOT ALLOW THE EXCLUSION OF IMPLIED WARRANTIES, SO SUCH EXCLUSION MAY NOT APPLY TO YOU.

6. Limitation on Liability. EXCEPT TO THE EXTENT REQUIRED BY APPLICABLE LAW, IN NO EVENT WILL LICENSOR BE LIABLE TO YOU ON ANY LEGAL THEORY FOR ANY SPECIAL, INCIDENTAL, CONSEQUENTIAL, PUNITIVE OR EXEMPLARY DAMAGES ARISING OUT OF THIS LICENSE OR THE USE OF THE WORK, EVEN IF LICENSOR HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

7. Termination

a. This License and the rights granted hereunder will terminate automatically upon any breach by You of the terms of this License. Individuals or entities who have received Derivative Works or Collective Works from You under this License, however, will not have their licenses terminated provided such individuals or entities remain in full compliance with those licenses. Sections 1, 2, 5, 6, 7, and 8 will survive any termination of this License.

b. Subject to the above terms and conditions, the license granted here is perpetual (for the duration of the applicable copyright in the Work). Notwithstanding the above, Licensor reserves the right to release the Work under different license terms or to stop distributing the Work at any time; provided, however that any such election will not serve to withdraw this License (or any other license that has been, or is required to be, granted under the terms of this License), and this License will continue in full force and effect unless terminated as stated above.

8. Miscellaneous

- a. Each time You distribute or publicly digitally perform the Work or a Collective Work, the Licensor offers to the recipient a license to the Work on the same terms and conditions as the license granted to You under this License.
- b. Each time You distribute or publicly digitally perform a Derivative Work, Licensor offers to the recipient a license to the original Work on the same terms and conditions as the license granted to You under this License.
- c. If any provision of this License is invalid or unenforceable under applicable law, it shall not affect the validity or enforceability of the remainder of the terms of this License, and without further action by the parties to this agreement, such provision shall be reformed to the minimum extent necessary to make such provision valid and enforceable.
- d. No term or provision of this License shall be deemed waived and no breach consented to unless such waiver or consent shall be in writing and signed by the party to be charged with such waiver or consent.
- e. This License constitutes the entire agreement between the parties with respect to the Work licensed here. There are no understandings, agreements or representations with respect to the Work not specified here. Licensor shall not be bound by any additional provisions that may appear in any communication from You. This License may not be modified without the mutual written agreement of the Licensor and You.

** Yahoo! Geoplanet; version latest -- <https://tinyurl.com/y8p87rwl>

Content licensed under Creative Commons Attribution 3.0 United States

License. Provide attribution to "Yahoo! GeoPlanet"

License

THE WORK (AS DEFINED BELOW) IS PROVIDED UNDER THE TERMS OF THIS CREATIVE COMMONS PUBLIC LICENSE ("CCPL" OR "LICENSE"). THE WORK IS PROTECTED BY COPYRIGHT AND/OR OTHER

APPLICABLE LAW. ANY USE OF THE WORK OTHER THAN AS AUTHORIZED UNDER THIS LICENSE OR COPYRIGHT LAW IS PROHIBITED.

BY EXERCISING ANY RIGHTS TO THE WORK PROVIDED HERE, YOU ACCEPT AND AGREE TO BE BOUND BY THE TERMS OF THIS LICENSE. TO THE EXTENT THIS LICENSE MAY BE CONSIDERED TO BE A CONTRACT, THE LICENSOR GRANTS YOU THE RIGHTS CONTAINED HERE IN CONSIDERATION OF YOUR ACCEPTANCE OF SUCH TERMS AND CONDITIONS.

1. Definitions

"Adaptation" means a work based upon the Work, or upon the Work and other pre-existing works, such as a translation, adaptation, derivative work, arrangement of music or other alterations of a literary or artistic work, or phonogram or performance and includes cinematographic adaptations or any other form in which the Work may be recast, transformed, or adapted including in any form recognizably derived from the original, except that a work that constitutes a Collection will not be considered an Adaptation for the purpose of this License. For the avoidance of doubt, where the Work is a musical work, performance or phonogram, the synchronization of the Work in timed-relation with a moving image ("synching") will be considered an Adaptation for the purpose of this License.

"Collection" means a collection of literary or artistic works, such as encyclopedias and anthologies, or performances, phonograms or broadcasts, or other works or subject matter other than works listed in Section 1(f) below, which, by reason of the selection and arrangement of their contents, constitute intellectual creations, in which the Work is included in its entirety in unmodified form along with one or more other contributions, each constituting separate and independent works in themselves, which together are assembled into a collective whole. A work that constitutes a Collection will not be considered an Adaptation (as defined above) for the purposes of this License.

"Distribute" means to make available to the public the original and copies of the Work or Adaptation, as appropriate, through sale or other transfer of ownership.

"Licensor" means the individual, individuals, entity or entities that offer(s) the Work under the terms of this License.

"Original Author" means, in the case of a literary or artistic work, the individual, individuals, entity or entities who created the Work or if no individual or entity can be identified, the publisher; and in addition (i) in the case of a performance the actors, singers, musicians, dancers, and other persons who act, sing, deliver, declaim, play in, interpret or otherwise perform literary or artistic

works or expressions of folklore; (ii) in the case of a phonogram the producer being the person or legal entity who first fixes the sounds of a performance or other sounds; and, (iii) in the case of broadcasts, the organization that transmits the broadcast.

"Work" means the literary and/or artistic work offered under the terms of this License including without limitation any production in the literary, scientific and artistic domain, whatever may be the mode or form of its expression including digital form, such as a book, pamphlet and other writing; a lecture, address, sermon or other work of the same nature; a dramatic or dramatico-musical work; a choreographic work or entertainment in dumb show; a musical composition with or without words; a cinematographic work to which are assimilated works expressed by a process analogous to cinematography; a work of drawing, painting, architecture, sculpture, engraving or lithography; a photographic work to which are assimilated works expressed by a process analogous to photography; a work of applied art; an illustration, map, plan, sketch or three-dimensional work relative to geography, topography, architecture or science; a performance; a broadcast; a phonogram; a compilation of data to the extent it is protected as a copyrightable work; or a work performed by a variety or circus performer to the extent it is not otherwise considered a literary or artistic work.

"You" means an individual or entity exercising rights under this License who has not previously violated the terms of this License with respect to the Work, or who has received express permission from the Licensor to exercise rights under this License despite a previous violation.

"Publicly Perform" means to perform public recitations of the Work and to communicate to the public those public recitations, by any means or process, including by wire or wireless means or public digital performances; to make available to the public Works in such a way that members of the public may access these Works from a place and at a place individually chosen by them; to perform the Work to the public by any means or process and the communication to the public of the performances of the Work, including by public digital performance; to broadcast and rebroadcast the Work by any means including signs, sounds or images.

"Reproduce" means to make copies of the Work by any means including without limitation by sound or visual recordings and the right of fixation and reproducing fixations of the Work, including storage of a protected performance or phonogram in digital form or other electronic medium.

2. Fair Dealing Rights. Nothing in this License is intended to reduce, limit, or restrict any uses free from copyright or rights arising from limitations or exceptions that are provided for in connection with the copyright protection under copyright law or other applicable laws.

3. License Grant. Subject to the terms and conditions of this License, Licensor hereby grants You a worldwide, royalty-free, non-exclusive, perpetual (for the duration of the applicable copyright) license to exercise the rights in the Work as stated below:

to Reproduce the Work, to incorporate the Work into one or more Collections, and to Reproduce the Work as incorporated in the Collections;

to create and Reproduce Adaptations provided that any such Adaptation, including any translation in any medium, takes reasonable steps to clearly label, demarcate or otherwise identify that changes were made to the original Work. For example, a translation could be marked "The original work was translated from English to Spanish," or a modification could indicate "The original work has been modified.";

to Distribute and Publicly Perform the Work including as incorporated in Collections; and,

to Distribute and Publicly Perform Adaptations.

For the avoidance of doubt:

Non-waivable Compulsory License Schemes. In those jurisdictions in which the right to collect royalties through any statutory or compulsory licensing scheme cannot be waived, the Licensor reserves the exclusive right to collect such royalties for any exercise by You of the rights granted under this License;

Waivable Compulsory License Schemes. In those jurisdictions in which the right to collect royalties through any statutory or compulsory licensing scheme can be waived, the Licensor waives the exclusive right to collect such royalties for any exercise by You of the rights granted under this License; and,

Voluntary License Schemes. The Licensor waives the right to collect royalties, whether individually or, in the event that the Licensor is a member of a collecting society that administers voluntary licensing schemes, via that society, from any exercise by You of the rights granted under this License.

The above rights may be exercised in all media and formats whether now known or hereafter devised. The above rights include the right to make such modifications as are technically necessary

to exercise the rights in other media and formats. Subject to Section 8(f), all rights not expressly granted by Licensor are hereby reserved.

4. Restrictions. The license granted in Section 3 above is expressly made subject to and limited by the following restrictions:

You may Distribute or Publicly Perform the Work only under the terms of this License. You must include a copy of, or the Uniform Resource Identifier (URI) for, this License with every copy of the Work You Distribute or Publicly Perform. You may not offer or impose any terms on the Work that restrict the terms of this License or the ability of the recipient of the Work to exercise the rights granted to that recipient under the terms of the License. You may not sublicense the Work. You must keep intact all notices that refer to this License and to the disclaimer of warranties with every copy of the Work You Distribute or Publicly Perform. When You Distribute or Publicly Perform the Work, You may not impose any effective technological measures on the Work that restrict the ability of a recipient of the Work from You to exercise the rights granted to that recipient under the terms of the License. This Section 4(a) applies to the Work as incorporated in a Collection, but this does not require the Collection apart from the Work itself to be made subject to the terms of this License. If You create a Collection, upon notice from any Licensor You must, to the extent practicable, remove from the Collection any credit as required by Section 4(b), as requested. If You create an Adaptation, upon notice from any Licensor You must, to the extent practicable, remove from the Adaptation any credit as required by Section 4(b), as requested.

If You Distribute, or Publicly Perform the Work or any Adaptations or Collections, You must, unless a request has been made pursuant to Section 4(a), keep intact all copyright notices for the Work and provide, reasonable to the medium or means You are utilizing: (i) the name of the Original Author (or pseudonym, if applicable) if supplied, and/or if the Original Author and/or Licensor designate another party or parties (e.g., a sponsor institute, publishing entity, journal) for attribution ("Attribution Parties") in Licensor's copyright notice, terms of service or by other reasonable means, the name of such party or parties; (ii) the title of the Work if supplied; (iii) to the extent reasonably practicable, the URI, if any, that Licensor specifies to be associated with the Work, unless such URI does not refer to the copyright notice or licensing information for the Work; and (iv) , consistent with Section 3(b), in the case of an Adaptation, a credit identifying the use of the Work in the Adaptation (e.g., "French translation of the Work by Original Author," or "Screenplay based on original Work by Original Author"). The credit required by this Section 4 (b) may be implemented in any reasonable manner; provided, however, that in the case of a Adaptation or Collection, at a minimum such credit will appear, if a credit for all contributing authors of the Adaptation or Collection appears, then as part of these credits and in a manner at

least as prominent as the credits for the other contributing authors. For the avoidance of doubt, You may only use the credit required by this Section for the purpose of attribution in the manner set out above and, by exercising Your rights under this License, You may not implicitly or explicitly assert or imply any connection with, sponsorship or endorsement by the Original Author, Licensor and/or Attribution Parties, as appropriate, of You or Your use of the Work, without the separate, express prior written permission of the Original Author, Licensor and/or Attribution Parties.

Except as otherwise agreed in writing by the Licensor or as may be otherwise permitted by applicable law, if You Reproduce, Distribute or Publicly Perform the Work either by itself or as part of any Adaptations or Collections, You must not distort, mutilate, modify or take other derogatory action in relation to the Work which would be prejudicial to the Original Author's honor or reputation. Licensor agrees that in those jurisdictions (e.g. Japan), in which any exercise of the right granted in Section 3(b) of this License (the right to make Adaptations) would be deemed to be a distortion, mutilation, modification or other derogatory action prejudicial to the Original Author's honor and reputation, the Licensor will waive or not assert, as appropriate, this Section, to the fullest extent permitted by the applicable national law, to enable You to reasonably exercise Your right under Section 3(b) of this License (right to make Adaptations) but not otherwise.

5. Representations, Warranties and Disclaimer

UNLESS OTHERWISE MUTUALLY AGREED TO BY THE PARTIES IN WRITING, LICENSOR OFFERS THE WORK AS-IS AND MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND CONCERNING THE WORK, EXPRESS, IMPLIED, STATUTORY OR OTHERWISE, INCLUDING, WITHOUT LIMITATION, WARRANTIES OF TITLE, MERCHANTIBILITY, FITNESS FOR A PARTICULAR PURPOSE, NONINFRINGEMENT, OR THE ABSENCE OF LATENT OR OTHER DEFECTS, ACCURACY, OR THE PRESENCE OF ABSENCE OF ERRORS, WHETHER OR NOT DISCOVERABLE. SOME JURISDICTIONS DO NOT ALLOW THE EXCLUSION OF IMPLIED WARRANTIES, SO SUCH EXCLUSION MAY NOT APPLY TO YOU.

6. Limitation on Liability. EXCEPT TO THE EXTENT REQUIRED BY APPLICABLE LAW, IN NO EVENT WILL LICENSOR BE LIABLE TO YOU ON ANY LEGAL THEORY FOR ANY SPECIAL, INCIDENTAL, CONSEQUENTIAL, PUNITIVE OR EXEMPLARY DAMAGES ARISING OUT OF THIS LICENSE OR THE USE OF THE WORK, EVEN IF LICENSOR HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

7. Termination

This License and the rights granted hereunder will terminate automatically upon any breach by You of the terms of this License. Individuals or entities who have received Adaptations or Collections from You under this License, however, will not have their licenses terminated provided such individuals or entities remain in full compliance with those licenses. Sections 1, 2, 5, 6, 7, and 8 will survive any termination of this License.

Subject to the above terms and conditions, the license granted here is perpetual (for the duration of the applicable copyright in the Work). Notwithstanding the above, Licensor reserves the right to release the Work under different license terms or to stop distributing the Work at any time; provided, however that any such election will not serve to withdraw this License (or any other license that has been, or is required to be, granted under the terms of this License), and this License will continue in full force and effect unless terminated as stated above.

8. Miscellaneous

Each time You Distribute or Publicly Perform the Work or a Collection, the Licensor offers to the recipient a license to the Work on the same terms and conditions as the license granted to You under this License.

Each time You Distribute or Publicly Perform an Adaptation, Licensor offers to the recipient a license to the original Work on the same terms and conditions as the license granted to You under this License.

If any provision of this License is invalid or unenforceable under applicable law, it shall not affect the validity or enforceability of the remainder of the terms of this License, and without further action by the parties to this agreement, such provision shall be reformed to the minimum extent necessary to make such provision valid and enforceable.

No term or provision of this License shall be deemed waived and no breach consented to unless such waiver or consent shall be in writing and signed by the party to be charged with such waiver or consent.

This License constitutes the entire agreement between the parties with respect to the Work licensed here. There are no understandings, agreements or representations with respect to the Work not specified here. Licensor shall not be bound by any additional provisions that may appear in any communication from You. This License may not be modified without the mutual written agreement of the Licensor and You.

The rights granted under, and the subject matter referenced, in this License were drafted utilizing the terminology of the Berne Convention for the Protection of Literary and Artistic Works (as amended on September 28, 1979), the Rome Convention of 1961, the WIPO Copyright Treaty of 1996, the WIPO Performances and Phonograms Treaty of 1996 and the Universal Copyright Convention (as revised on July 24, 1971). These rights and subject matter take effect in the relevant jurisdiction in which the License terms are sought to be enforced according to the corresponding provisions of the implementation of those treaty provisions in the applicable national law. If the standard suite of rights granted under applicable copyright law includes additional rights not granted under this License, such additional rights are deemed to be included in the License; this License is not intended to restrict the license of any rights under applicable law.

** geonames.org; version latest -- <https://www.geonames.org>

This work is licensed under a Creative Commons Attribution 4.0 License, see <http://creativecommons.org/licenses/by/4.0/>

The Data is provided "as is" without warranty or any representation of accuracy, timeliness or completeness.

** Zetashapes; version latest -- <https://tinyurl.com/y8p87rwl>

The basic source data is from US TIGER/Line Census Data which is public domain (Q10). This site also makes use of data scraped from the flickr api.

Creative Commons Attribution 4.0 International Public License

By exercising the Licensed Rights (defined below), You accept and agree to be bound by the terms and conditions of this Creative Commons Attribution 4.0 International Public License ("Public License"). To the extent this Public License may be interpreted as a contract, You are granted the Licensed Rights in consideration of Your acceptance of these terms and conditions, and the Licensor grants You such rights in consideration of benefits the Licensor receives from making the Licensed Material available under these terms and conditions.

Section 1 – Definitions.

a. Adapted Material means material subject to Copyright and Similar Rights that is derived from or based upon the Licensed Material and in which the Licensed Material is translated, altered,

arranged, transformed, or otherwise modified in a manner requiring permission under the Copyright and Similar Rights held by the Licensor. For purposes of this Public License, where the Licensed Material is a musical work, performance, or sound recording, Adapted Material is always produced where the Licensed Material is synched in timed relation with a moving image.

b. Adapter's License means the license You apply to Your Copyright and Similar Rights in Your contributions to Adapted Material in accordance with the terms and conditions of this Public License.

c. Copyright and Similar Rights means copyright and/or similar rights closely related to copyright including, without limitation, performance, broadcast, sound recording, and Sui Generis Database Rights, without regard to how the rights are labeled or categorized. For purposes of this Public License, the rights specified in Section 2(b)(1)-(2) are not Copyright and Similar Rights.

d. Effective Technological Measures means those measures that, in the absence of proper authority, may not be circumvented under laws fulfilling obligations under Article 11 of the WIPO Copyright Treaty adopted on December 20, 1996, and/or similar international agreements.

e. Exceptions and Limitations means fair use, fair dealing, and/or any other exception or limitation to Copyright and Similar Rights that applies to Your use of the Licensed Material.

f. Licensed Material means the artistic or literary work, database, or other material to which the Licensor applied this Public License.

g. Licensed Rights means the rights granted to You subject to the terms and conditions of this Public License, which are limited to all Copyright and Similar Rights that apply to Your use of the Licensed Material and that the Licensor has authority to license.

h. Licensor means the individual(s) or entity(ies) granting rights under this Public License.

i. Share means to provide material to the public by any means or process that requires permission under the Licensed Rights, such as reproduction, public display, public performance, distribution, dissemination, communication, or importation, and to make material available to the public including in ways that members of the public may access the material from a place and at a time individually chosen by them.

j. Sui Generis Database Rights means rights other than copyright resulting from Directive 96/9/EC of the European Parliament and of the Council of 11 March 1996 on the legal protection of databases, as amended and/or succeeded, as well as other essentially equivalent rights anywhere in the world.

k. You means the individual or entity exercising the Licensed Rights under this Public License. Your has a corresponding meaning.

Section 2 – Scope.

a. License grant.

1. Subject to the terms and conditions of this Public License, the Licensor hereby grants You a worldwide, royalty-free, non-sublicensable, non-exclusive, irrevocable license to exercise the Licensed Rights in the Licensed Material to:

A. reproduce and Share the Licensed Material, in whole or in part; and

B. produce, reproduce, and Share Adapted Material.

2. Exceptions and Limitations. For the avoidance of doubt, where Exceptions and Limitations apply to Your use, this Public License does not apply, and You do not need to comply with its terms and conditions.

3. Term. The term of this Public License is specified in Section 6(a).

4. Media and formats; technical modifications allowed. The Licensor authorizes You to exercise the Licensed Rights in all media and formats whether now known or hereafter created, and to make technical modifications necessary to do so. The Licensor waives and/or agrees not to assert any right or authority to forbid You from making technical modifications necessary to exercise the Licensed Rights, including technical modifications necessary to circumvent Effective Technological Measures. For purposes of this Public License, simply making modifications authorized by this Section 2(a)(4) never produces Adapted Material.

5. Downstream recipients.

A. Offer from the Licensor – Licensed Material. Every recipient of the Licensed Material automatically receives an offer from the Licensor to exercise the Licensed Rights under the terms and conditions of this Public License.

B. No downstream restrictions. You may not offer or impose any additional or different terms or conditions on, or apply any Effective Technological Measures to, the Licensed Material if doing so restricts exercise of the Licensed Rights by any recipient of the Licensed Material.

6. No endorsement. Nothing in this Public License constitutes or may be construed as permission to assert or imply that You are, or that Your use of the Licensed Material is, connected with, or

sponsored, endorsed, or granted official status by, the Licensor or others designated to receive attribution as provided in Section 3(a)(1)(A)(i).

b. Other rights.

1. Moral rights, such as the right of integrity, are not licensed under this Public License, nor are publicity, privacy, and/or other similar personality rights; however, to the extent possible, the Licensor waives and/or agrees not to assert any such rights held by the Licensor to the limited extent necessary to allow You to exercise the Licensed Rights, but not otherwise.

2. Patent and trademark rights are not licensed under this Public License.

3. To the extent possible, the Licensor waives any right to collect royalties from You for the exercise of the Licensed Rights, whether directly or through a collecting society under any voluntary or waivable statutory or compulsory licensing scheme. In all other cases the Licensor expressly reserves any right to collect such royalties.

Section 3 – License Conditions.

Your exercise of the Licensed Rights is expressly made subject to the following conditions.

a. Attribution.

1. If You Share the Licensed Material (including in modified form), You must:

A. retain the following if it is supplied by the Licensor with the Licensed Material:

i. identification of the creator(s) of the Licensed Material and any others designated to receive attribution, in any reasonable manner requested by the Licensor (including by pseudonym if designated);

ii. a copyright notice;

iii. a notice that refers to this Public License;

iv. a notice that refers to the disclaimer of warranties;

v. a URI or hyperlink to the Licensed Material to the extent reasonably practicable;

B. indicate if You modified the Licensed Material and retain an indication of any previous modifications; and

C. indicate the Licensed Material is licensed under this Public License, and include the text of, or the URI or hyperlink to, this Public License.

2. You may satisfy the conditions in Section 3(a)(1) in any reasonable manner based on the medium, means, and context in which You Share the Licensed Material. For example, it may be reasonable to satisfy the conditions by providing a URI or hyperlink to a resource that includes the required information.

3. If requested by the Licensor, You must remove any of the information required by Section 3(a)(1) (A) to the extent reasonably practicable.

4. If You Share Adapted Material You produce, the Adapter's License You apply must not prevent recipients of the Adapted Material from complying with this Public License.

Section 4 – Sui Generis Database Rights.

Where the Licensed Rights include Sui Generis Database Rights that apply to Your use of the Licensed Material:

- a. for the avoidance of doubt, Section 2(a)(1) grants You the right to extract, reuse, reproduce, and Share all or a substantial portion of the contents of the database;
- b. if You include all or a substantial portion of the database contents in a database in which You have Sui Generis Database Rights, then the database in which You have Sui Generis Database Rights (but not its individual contents) is Adapted Material; and
- c. You must comply with the conditions in Section 3(a) if You Share all or a substantial portion of the contents of the database.

For the avoidance of doubt, this Section 4 supplements and does not replace Your obligations under this Public License where the Licensed Rights include other Copyright and Similar Rights.

Section 5 – Disclaimer of Warranties and Limitation of Liability.

a. Unless otherwise separately undertaken by the Licensor, to the extent possible, the Licensor offers the Licensed Material as-is and as-available, and makes no representations or warranties of any kind concerning the Licensed Material, whether express, implied, statutory, or other. This includes, without limitation, warranties of title, merchantability, fitness for a particular purpose, non-infringement, absence of latent or other defects, accuracy, or the presence or absence of errors, whether or not known or discoverable. Where disclaimers of warranties are not allowed in full or in part, this disclaimer may not apply to You.

b. To the extent possible, in no event will the Licensor be liable to You on any legal theory (including, without limitation, negligence) or otherwise for any direct, special, indirect, incidental, consequential, punitive, exemplary, or other losses, costs, expenses, or damages arising out of this Public License or use of the Licensed Material, even if the Licensor has been advised of the possibility of such losses, costs, expenses, or damages. Where a limitation of liability is not allowed in full or in part, this limitation may not apply to You.

c. The disclaimer of warranties and limitation of liability provided above shall be interpreted in a manner that, to the extent possible, most closely approximates an absolute disclaimer and waiver of all liability.

Section 6 – Term and Termination.

a. This Public License applies for the term of the Copyright and Similar Rights licensed here. However, if You fail to comply with this Public License, then Your rights under this Public License terminate automatically.

b. Where Your right to use the Licensed Material has terminated under Section 6(a), it reinstates:

1. automatically as of the date the violation is cured, provided it is cured within 30 days of Your discovery of the violation; or
2. upon express reinstatement by the Licensor.

For the avoidance of doubt, this Section 6(b) does not affect any right the Licensor may have to seek remedies for Your violations of this Public License.

c. For the avoidance of doubt, the Licensor may also offer the Licensed Material under separate terms or conditions or stop distributing the Licensed Material at any time; however, doing so will not terminate this Public License.

d. Sections 1, 5, 6, 7, and 8 survive termination of this Public License.

Section 7 – Other Terms and Conditions.

a. The Licensor shall not be bound by any additional or different terms or conditions communicated by You unless expressly agreed.

b. Any arrangements, understandings, or agreements regarding the Licensed Material not stated herein are separate from and independent of the terms and conditions of this Public License.

Section 8 – Interpretation.

- a. For the avoidance of doubt, this Public License does not, and shall not be interpreted to, reduce, limit, restrict, or impose conditions on any use of the Licensed Material that could lawfully be made without permission under this Public License.
- b. To the extent possible, if any provision of this Public License is deemed unenforceable, it shall be automatically reformed to the minimum extent necessary to make it enforceable. If the provision cannot be reformed, it shall be severed from this Public License without affecting the enforceability of the remaining terms and conditions.
- c. No term or condition of this Public License will be waived and no failure to comply consented to unless expressly agreed to by the Licensor.
- d. Nothing in this Public License constitutes or may be interpreted as a limitation upon, or waiver of, any privileges and immunities that apply to the Licensor or You, including from the legal processes of any jurisdiction or authority.

** Microsoft Sample Databases: World Wide Importers; version latest -- <https://tinyurl.com/ybat5d8k>

Microsoft SQL Server Sample Code

Copyright (c) Microsoft Corporation

All rights reserved.

MIT License.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

**** Microsoft Sample Databases: AdventureWorks; version last -- <https://github.com/Microsoft/sql-server-samples/releases/tag/adventureworks2014>**

Microsoft SQL Server Sample Code

Copyright (c) Microsoft Corporation

All rights reserved.

Microsoft Public License (Ms-PL)

This license governs use of the accompanying software. If you use the software, you accept this license. If you do not accept the license, do not use the software.

1. Definitions

The terms "reproduce," "reproduction," "derivative works," and "distribution" have the same meaning here as under U.S. copyright law.

A "contribution" is the original software, or any additions or changes to the software.

A "contributor" is any person that distributes its contribution under this license.

"Licensed patents" are a contributor's patent claims that read directly on its contribution.

2. Grant of Rights

(A) Copyright Grant- Subject to the terms of this license, including the license conditions and limitations in section 3, each contributor grants you a non-exclusive, worldwide, royalty-free copyright license to reproduce its contribution, prepare derivative works of its contribution, and distribute its contribution or any derivative works that you create.

(B) Patent Grant- Subject to the terms of this license, including the license conditions and limitations in section 3, each contributor grants you a non-exclusive, worldwide, royalty-free license under its licensed patents to make, have made, use, sell, offer for sale, import, and/or otherwise dispose of its contribution in the software or derivative works of the contribution in the software.

3. Conditions and Limitations

(A) No Trademark License- This license does not grant you rights to use any contributors' name, logo, or trademarks.

(B) If you bring a patent claim against any contributor over patents that you claim are infringed by the software, your patent license from such contributor to the software ends automatically.

(C) If you distribute any portion of the software, you must retain all copyright, patent, trademark, and attribution notices that are present in the software.

(D) If you distribute any portion of the software in source code form, you may do so only under this license by including a complete copy of this license with your distribution. If you distribute any portion of the software in compiled or object code form, you may only do so under a license that complies with this license.

(E) The software is licensed "as-is." You bear the risk of using it. The contributors give no express warranties, guarantees or conditions. You may have additional consumer rights under your local laws which this license cannot change. To the extent permitted under your local laws, the contributors exclude the implied warranties of merchantability, fitness for a particular purpose and non-infringement.

** OurAirports; version latest -- <https://tinyurl.com/y8p87rwl>

Made with OurAirports. <http://ourairports.com/data/>

We hereby release all of these files into the Public Domain, with no warranty of any kind — By downloading any of these files, you agree that OurAirports.com, Megginson Technologies Ltd., and anyone involved with the web site or company hold no liability for anything that happens when you use the data, including (but not limited to) computer damage, lost revenue, flying into cliffs, or a general feeling of drowsiness that persists more than two days.

** Natural Earth; version latest -- <https://tinyurl.com/y8p87rwl>

Made with Natural Earth. Free vector and raster map data @ naturalearthdata.com.

All versions of Natural Earth raster + vector map data found on this

website are in the public domain. You may use the maps in any manner, including modifying the content and design, electronic dissemination, and offset printing. The primary authors, Tom Patterson and Nathaniel Vaughn Kelso, and all other contributors renounce all financial claim to the maps and invites you to use them for personal, educational, and commercial purposes.

No permission is needed to use Natural Earth. Crediting the authors is unnecessary.